

---

---

PLEASE NOTE:

The following selection from *The Yale Law Journal* is not paginated and should not be used for citation purposes. A paginated version of the selection may be purchased from *The Yale Law Journal*.

Notwithstanding any designation to the contrary in any printed version, copyright is held by the author.

© 2002 Yochai Benkler

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License. To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/1.0>

or send a letter to Creative Commons,  
559 Nathan Abbott Way, Stanford,  
California 94305, USA.

## Article

# Coase's Penguin, or, Linux and *The Nature of the Firm*

Yochai Benkler<sup>†</sup>

## CONTENTS

INTRODUCTION.....
I. PEER PRODUCTION ALL AROUND .....
A. Content.....
B. Relevance/Accreditation .....
C. Value-Added Distribution .....
D. Summary .....
II. WHY WOULD PEER PRODUCTION EMERGE IN A NETWORKED ENVIRONMENT?.....
A. Locating the Theoretical Space for Peer Production .....

---

<sup>†</sup> Professor of Law, New York University School of Law. Research for this Article was partly supported by a grant from the Filomen D'Agostino and Max E. Greenberg Research Fund at the New York University School of Law. I owe thanks to many for comments on this and earlier drafts, including: Bruce Ackerman, Ed Baker, Elazar Barkan, Jamie Boyle, Dan Burk, Ben Damstedt, Bob Ellickson, Terry Fisher, Natalie Jeremijenko, David Johnson, Dan Kahan, Niva Elkin Koren, Tara Lemmey, Larry Lessig, Doug Lichtman, Jerry Mashaw, Eben Moglen, Mark Nadel, Helen Nissenbaum, George Priest, Peggy Radin, Carol Rose, Chuck Sabel, Alan Schwartz, Clay Shirky, Richard Stallman, and Kenji Yoshino. I owe special thanks to Steve Snyder for his invaluable research assistance on the peer production enterprises described here.

I have gotten many questions about the “Coase’s Penguin” portion of the title. It turns out that the geek culture that easily recognizes “Coase” doesn’t recognize the “Penguin,” and vice versa. “Coase” refers to Ronald Coase, who originated the transaction costs theory of the firm that provides the methodological template for the positive analysis of peer production that I offer here. The penguin refers to the fact that the Linux kernel development community has adopted the image of a paunchy penguin as its mascot/trademark. One result of this cross-cultural conversation is that I will occasionally explain in some detail concepts that are well known in one community but not in the other.

2002]

Coase's Penguin

B.	<i>Peer Production of Information in a Pervasively Networked Environment</i> .....
C.	<i>Markets, Hierarchies, and Peer Production as Information-Processing Systems</i> .....
1.	<i>Information Gains</i> .....
2.	<i>Allocation Gains</i> .....
III. OF MOTIVATION AND ORGANIZATION: THE COMMONS PROBLEM .....	
A.	<i>The "Incentives" Problem: Of Diverse Motivations and Small Contributions</i> .....
1.	<i>Abstracting the Effect of Diverse Motivations</i> .....
2.	<i>Diverse Motivations and Large-Scale Collaborations</i> .....
B.	<i>Integration: Problem and Opportunity</i> .....
1.	<i>Threats to Motivation</i> .....
2.	<i>Provisioning Integration</i> .....
CONCLUSION .....	

## INTRODUCTION

Imagine that back in the days when what was good for GM was good for the country, an advisory committee of economists had recommended to the President of the United States that the federal government should support the efforts of volunteer communities to design and build cars, either for sale or for free distribution to automobile drivers. The committee members would probably have been locked up in a psychiatric ward—if Senator McCarthy or the House Un-American Activities Committee did not get them first. Yet, in September 2000, something like this actually happened. The President's Information Technology Advisory Committee recommended that the federal government support open source software as a strategic national choice to sustain the U.S. lead in critical software development.<sup>1</sup>

At the heart of the economic engine of the world's most advanced economies, and in particular that of the United States, we are beginning to take notice of a hardy, persistent, and quite amazing phenomenon. A new model of production has taken root, one that should not be there, at least

---

1. PRESIDENT'S INFO. TECH. ADVISORY COMM., DEVELOPING OPEN SOURCE SOFTWARE TO ADVANCE HIGH END COMPUTING (2000), *at* <http://www.ccic.gov/pubs/pitac/pres-oss-11sep00.pdf>.

according to our most widely held beliefs about economic behavior. The intuitions of the late twentieth-century American resist the idea that thousands of volunteers could collaborate on a complex economic project. It certainly should not be that these volunteers will beat the largest and best-financed business enterprises in the world at their own game. And yet, this is precisely what is happening in the software industry.

The emergence of free software<sup>2</sup> and the phenomenal success of its flagships—the GNU/Linux operating system,<sup>3</sup> the Apache web server, Perl, sendmail, BIND—and many other projects<sup>4</sup> should force us to take a second look at the dominant paradigm we hold about productivity. In the late 1930s, Ronald Coase wrote *The Nature of the Firm*,<sup>5</sup> in which he explained why firms emerge, defining firms as clusters of resources and agents that interact through managerial command systems rather than markets. In that paper, Coase introduced the concept of transaction costs, which are costs associated with defining and enforcing property and contract rights and which are a necessary incident of organizing any activity on a market model. Coase explained the emergence and limits of firms based on the differences in the transaction costs associated with organizing production through markets or through firms. People use markets when the gains from doing so, net of transaction costs, exceed the gains from doing the same thing in a managed firm, net of organization costs. Firms emerge when the opposite is true. Any individual firm will stop growing when its organization costs exceed the organization costs of a smaller firm. This basic insight was then extended and developed in the work of Oliver

---

2. I use the terms “free software” and “open source software” interchangeably in this Article. Those who consider the phenomenon as first and foremost involving political values, to wit, freedom, use the former, in self-conscious contradistinction to those who focus on the economic significance, who use the latter. Compare ERIC RAYMOND, *Homesteading the Noosphere, in THE CATHEDRAL AND THE BAZAAR: MUSINGS ON LINUX AND OPEN SOURCE BY AN ACCIDENTAL REVOLUTIONARY* 65 (2001) [hereinafter THE CATHEDRAL AND THE BAZAAR] (focusing on the economic significance), available at [http://www.firstmonday.dk/issues/issue3\\_10/raymond/](http://www.firstmonday.dk/issues/issue3_10/raymond/), with Free Software Found., Why “Free Software” Is Better than “Open Source,” at <http://www.fsf.org/philosophy/free-software-for-freedom.html> (last modified Aug. 26, 2002) (focusing on the political significance). I have written and continue to write quite extensively on the normative implications of how information production is organized, see, e.g., Yochai Benkler, *The Battle over the Institutional Ecosystem in the Digital Environment*, 44 COMM. ACM 84 (2001), but not in this Article, where I generally abjure disputes over the terms.

3. I describe the operating system as GNU/Linux to denote that it is a combination of the kernel development project initiated by Linus Torvalds in 1991 and of many other operating system components created by the GNU project, which was originated in 1984 by Richard Stallman, the father of free software. Throughout the Article, I refer to GNU or Linux separately to denote the specific development project and to the operating system as GNU/Linux. I departed from this practice in the title for stylistic purposes alone. The complete GNU/Linux operating system is what everyone has in mind when they speak of the breathtaking success of free software at making excellent high-end software.

4. For an excellent history of the free software movement and of the open source development methodology, see GLYN MOODY, *REBEL CODE* (2001).

5. Ronald H. Coase, *The Nature of the Firm*, 4 ECONOMICA 386 (1937).

2002]

## Coase's Penguin

Williamson and other institutional economists who studied the relationship between markets and managerial hierarchies as models of organizing production.<sup>6</sup>

The emergence of free software as a substantial force in the software-development world poses a puzzle for this organization theory. Free software projects do not rely either on markets or on managerial hierarchies to organize production. Programmers do not generally participate in a project because someone who is their boss instructed them, though some do. They do not generally participate in a project because someone offers them a price, though some participants do focus on long-term appropriation through money-oriented activities, like consulting or service contracts. But the critical mass of participation in projects cannot be explained by the direct presence of a command, a price, or even a future monetary return, particularly in the all-important microlevel decisions regarding selection of projects to which participants contribute.<sup>7</sup> In other words, programmers participate in free software projects without following the normal signals generated by market-based, firm-based, or hybrid models.

This puzzle has attracted increasing attention from economists<sup>8</sup> and participants in the practice<sup>9</sup> trying to understand their own success and its

---

6. The initial framing in terms of the opposition between markets and hierarchy was OLIVER E. WILLIAMSON, MARKETS AND HIERARCHIES: ANALYSIS AND ANTITRUST IMPLICATIONS: A STUDY IN THE ECONOMICS OF INTERNAL ORGANIZATION (1975) and OLIVER E. WILLIAMSON, THE ECONOMIC INSTITUTIONS OF CAPITALISM (1985). See also Benjamin Klein et al., *Vertical Integration, Appropriable Rents, and the Competitive Contracting Process*, 21 J.L. & ECON. 297 (1978) (discussing contractual relationships as blurring the line between markets and firms). State hierarchies are also an option, and while the extreme version—socialist production—is largely discredited, some state production of some goods, like power, is still very much in play. Here, I focus only on market production, whether decentralized and price-driven or firm-based and managed. Any arguments about the importance of governmental investment in science, research, and the arts are independent of the potential conclusions for intellectual property that this Article suggests.

7. Even if it could be established, as it has not, that most contributors to free software development projects were motivated by extrinsic monetary rewards, like gaining consulting contracts through reputation and human capital gains, price would still be of small explanatory value if those motivations led to a general willingness to contribute to *some* project but did not direct the actual selection of projects and type of contribution. It is revealing that while reputation is perhaps the most readily available and widely cited extrinsic motivator to contribution, its explanatory force wanes when the practices of two of the most successful free software projects are considered. Neither the Apache project nor the Free Software Foundation publishes the names of individual contributors. It is possible that reputation creation and flow is a more complex social phenomenon within the high priesthood than would be implied by explicit attribution, or that the star status of the highest priests is a sufficient reputation-based reward. It is also possible—indeed likely—that people's motivations are heterogeneous and that some people are more driven by explicit reputation gains than others. Whether people who are more driven by explicit reputation rewards will indeed cluster in projects where explicit reputation rewards are better organized remains a question that has not yet been studied empirically.

8. An excellent overview of, and insightful contribution to, this literature is the working paper STEVEN WEBER, THE POLITICAL ECONOMY OF OPEN SOURCE (Berkeley Roundtable on the Int'l Econ., Working Paper No. 140, 2000), at <http://brie.berkeley.edu/~briewww/pubs/wp/wp140.pdf>.

9. The canonical references here are to two works by Eric Raymond, an open source software developer who turned into the most vocal and widely read commentator on the phenomenon. ERIC

sustainability given widespread contrary intuitions. Josh Lerner and Jean Tirole present the best overarching view of the range of diverse micromotivations that drive free software developers.<sup>10</sup> This diversity of motivations, somewhat more formalized and generalized, plays an important role in my own analysis. Some writing by both practitioners and observers, supporters and critics, has focused on the “hacker ethic,” and analogized the sociological phenomenon to gift-exchange systems.<sup>11</sup> Other writing has focused on the special characteristics of software as an object of production.<sup>12</sup>

In this Article, I approach this puzzle by departing from free software. Rather than trying to explain what is special about software or hackers, I generalize from the phenomenon of free software to suggest characteristics that make large-scale collaborations in many information production fields sustainable and productive in the digitally networked environment without reliance either on markets or managerial hierarchy.<sup>13</sup> Hence the title of this

---

RAYMOND, *The Cathedral and the Bazaar*, in THE CATHEDRAL AND THE BAZAAR, *supra* note 2, at 19, available at [http://www.firstmonday.dk/issues/issue3\\_3/raymond/](http://www.firstmonday.dk/issues/issue3_3/raymond/); RAYMOND, *supra* note 2.

10. See Josh Lerner & Jean Tirole, *Some Simple Economics of Open Source*, 50 J. INDUS. ECON. 197, 212-23 (2002). Eric von Hippel has provided both theoretical and empirical support for the importance of the use value gained by users in a user-driven innovation environment, both in software and elsewhere. See Eric von Hippel, *Innovation by User Communities: Learning from Open-Source Software*, 42 SLOAN MGMT. REV. 82 (2001); Eric von Hippel, at <http://web.mit.edu/evhippel/www/Publications.htm> (last visited Aug. 25, 2002) (providing access to many collaborative papers); see also Jean-Michel Dalle & Nicolas Jullien, “Libre” Software: Turning Fads into Institutions (Jan. 18, 2001), at <http://opensource.mit.edu/papers/Libre-Software.pdf> (analyzing the heterogeneous motivations of contributors in free software projects).

11. In addition to Raymond, supporters of the sustainability of free software development who have used this framework include Rishab Aiyer Ghosh, *Cooking Pot Markets: An Economic Model for the Trade in Free Goods and Services on the Internet*, 3 FIRST MONDAY 3 (Mar. 2, 1998), at [http://www.firstmonday.dk/issues/issue3\\_3/ghosh/](http://www.firstmonday.dk/issues/issue3_3/ghosh/). See also Peter Kollock, *The Economies of Online Cooperation: Gifts and Public Goods in Cyberspace*, in COMMUNITIES IN CYBERSPACE 220 (Marc A. Smith & Peter Kollock eds., 1999). Less sanguine views of this development model, which are also based on the hacker-ethic framework, include Robert L. Glass, *The Sociology of Open Source: Of Cults and Cultures*, IEEE SOFTWARE, May-June 2000, at 104, and David Lancashire, *Code, Culture and Cash: The Fading Altruism of Open Source Development*, 6 FIRST MONDAY 3 (Dec. 3, 2001), at [http://www.firstmonday.org/issues/issue6\\_12/lancashire/](http://www.firstmonday.org/issues/issue6_12/lancashire/). For a discussion of the hacker ethic generally, not solely in the context of free software development, see PEKKA HIMANEN, LINUS TORVALDS & MANUEL CASTELLS, THE HACKER ETHIC (2001).

12. See, e.g., James Bessen, Open Source Software: Free Provision of Complex Public Goods (July 2002), at <http://www.researchoninnovation.org/opensrc.pdf>.

13. The most closely related work in the open source software literature is the mapping of diverse motivations, *see supra* notes 10-11, and those papers that try to explain the open source software development model in terms of its information-sharing characteristics, *see, e.g.*, Justin Pappas Johnson, Economics of Open Source Software (May 17, 2001), at <http://opensource.mit.edu/papers/johnsonopensource.pdf> (recognizing superior access to the talent pool, but cautioning that free-riding will lead to underutilization); Anca Metiu & Bruce Kogut, Distributed Knowledge and the Global Organization of Software Development (Feb. 2001), at <http://opensource.mit.edu/papers/kogut1.pdf> (claiming that the value of a globally distributed skill set will loosen the grip of the richest countries on innovation).

2002]

## Coase's Penguin

Article—to invoke the challenge that the paunchy penguin mascot of the Linux kernel development community poses for the view of organization rooted in Coase's work.

Part I begins to tell the tale of the more general phenomenon through a number of detailed stories. Tens of thousands of individuals collaborate in five-minute increments to map Mars's craters, fulfilling tasks that would normally be performed by full-time Ph.D.s. A quarter of a million people collaborate on creating the most important news and commentary site currently available on technology issues. Twenty-five thousand people collaborate to create a peer-reviewed publication of commentary on technology and culture. Forty thousand people collaborate to create a more efficient human-edited directory for the Web than Yahoo. I offer other examples as well.

The point of Part I is simple. The phenomenon of large- and medium-scale collaborations among individuals that are organized without markets or managerial hierarchies is emerging everywhere in the information and cultural production system. The question is how we should understand these instances of socially productive behavior: What are the dynamics that make them possible and successful, and how should we think about their economic value?

My basic framework for explaining this emerging phenomenon occupies Part II of the Article. Collaborative production systems pose an information problem. The question that individual agents in such a system need to solve in order to be productive is what they should do. Markets solve this problem by attaching price signals to alternative courses of action. Firms solve this problem by assigning different signals from different agents different weights. To wit, what a manager says matters. In order to perform these functions, both markets and firms need to specify the object of the signal sufficiently so that property, contract, and managerial instructions can be used to differentiate between agents, efforts, resources,

---

The only treatment that specifically uses aspects of Coase's *The Nature of the Firm* as an analytic framework for understanding free software is David McGowan, *Legal Implications of Open-Source Software*, 2001 U. ILL. L. REV. 241. Congruent with Coase's conclusion, McGowan assumes that hierarchical control is necessary to coordinate projects in the absence of markets, and he demonstrates this effect as applied to the Linux kernel development process. He then analyzes how the licensing provisions and the social motivations and relationships involved in open source software projects form the basis for the hierarchical aspects of this software development model. *Id.* at 275-88. My own use of Coase's insights is very different. See *infra* Part II. I apply Coase's insight regarding the centrality of comparative transaction costs to the organizational form that a production enterprise will take. In my model, "information opportunity costs" play a similar role in describing the comparative social cost of different organizational forms to the role played by transaction costs more generally in the Coasean framework. Peer production emerges, as firms do in Coase's analysis, because it can have lower information opportunity costs under certain technological and economic conditions. McGowan's analysis therefore primarily intersects with this Article where I suggest that the integration in peer production processes sometimes takes the form of a hierarchy.

and potential combinations thereof. Where agents, efforts, or resources cannot be so specified, they cannot be accurately priced or managed. The process of specification creates two sources of inefficiency. First, it causes information loss. Perfect specification is unattainable because of transaction costs associated with specifying the characteristics of each human and material resource and each opportunity for utilization. Second, property and contract make clusters of agents and resources sticky. A firm's employees will more readily work with a firm's owned resources than with other sources and will more readily collaborate with other employees of the firm than with outsiders. It is not impossible to acquire and trade resources and collaborative efforts, but this is done only when the perceived gains outweigh the transaction costs. Nonproprietary production strategies can improve on markets and firms by correcting these two failures.

Commons-based peer production, the emerging third model of production I describe here, relies on decentralized information gathering and exchange to reduce the uncertainty of participants. It has particular advantages as an information process for identifying and allocating human creativity available to work on information and cultural resources.<sup>14</sup> It depends on very large aggregations of individuals independently scouring their information environment in search of opportunities to be creative in small or large increments. These individuals then self-identify for tasks and perform them for a variety of motivational reasons that I discuss at some length.

If the problems of motivation and organization can be solved, then commons-based peer production has two major advantages over firms and markets. First, it places the point of decision about assigning any given person to any given set of resources with the individual. Given the high variability among individuals and across time in terms of talent, experience, motivation, focus, availability, etc., human creativity is an especially difficult resource to specify for efficient contracting or management. Firms recognize this and attempt to solve this problem by creating various incentive compensation schemes and intangible reward schemes, like employee-of-the-month awards. These schemes work to some extent to alleviate the information loss associated with managerial production, but

---

14. This third mode of production is in some measure similar to the artisan mode of production identified by the path-breaking work MICHAEL J. PIORE & CHARLES F. SABLE, THE SECOND INDUSTRIAL DIVIDE (1984). There are, however, sufficient qualitative differences that make this a new phenomenon requiring its own set of understandings, rather than a latter-day artisan cooperative. Most important are the scale of these collaborations, the absence of entry barriers in many or most of them, and the absence of direct appropriation of the products. With regard to organization literature, commons-based peer production stands in a similar relationship to artisan production as, in the property literature, commons relate to common property regimes. These are phenomena that share several characteristics, but ultimately diverge in central characteristics that require different explanations.

2002]

Coase's Penguin

only insofar as a firm's agents and resources are indeed the best and only insofar as these schemes capture all the motivations and contributions accurately. Peer production provides a framework within which individuals who have the best information available about their own fit for a task can self-identify for the task. This provides an information gain over firms and markets, but only if the system develops some mechanism to filter out mistaken judgments that agents make about themselves. This is why practically all successful peer production systems have a robust mechanism for peer review or statistical weeding out of contributions from agents who misjudge themselves.

The allocation gains of peer production are as important as the information gains. Human creativity cannot be assumed to be an on-off switch of suitability for a job, as simple models of industrial production might treat labor. One cannot say in the information context that "this person passes threshold suitability requirements to pull this lever all day" and ignore variability beyond that fact. It is more likely that variability in productivity will be large for different people with any given set of resources and collaborators for any given set of projects. I describe this diversity as a probability that any agent will be a good fit with a set of resources and agents to produce valuable new information or cultural goods. Peer production has an advantage over firms and markets because it allows larger groups of individuals to scour larger groups of resources in search of materials, projects, collaborations, and combinations than is possible for firms or individuals who function in markets. Transaction costs associated with property and contract limit the access of people to each other, to resources, and to projects when production is organized on a market or firm model, but not when it is organized on a peer production model.<sup>15</sup> Because fit of people to projects and to each other is variable, there are increasing returns to the scale of the number of people, resources, and projects capable of being combined.

The advantages of peer production are, then, improved identification and allocation of human creativity. These advantages appear to have become salient, because human creativity itself has become salient. In the domain of information and culture, production generally comprises the combination of preexisting information/cultural inputs, human creativity, and the physical capital necessary to (1) fix ideas and human utterances in media capable of storing and communicating them and (2) transmit them.

---

15. This is not to say that there are no transaction costs associated with peer production. It is merely to say that these transaction costs, which largely fall under the rubric of "integration" that I describe in Section III.B, are of a different type. They may undermine the successful integration of a project or may make participation too costly for contributors, but they do not arise as a barrier to prevent many individuals from collaborating in the same resource space or to prevent many resources from populating that space.

Existing information and culture are a public good in the strict economic sense of being nonrival.<sup>16</sup> The cost of physical capital was for more than 150 years the central organizing principle of information and cultural production, from the introduction of high-cost, high-volume mechanical presses, through telegraph, telephone, radio, film, records, television, cable, and satellite systems. These costs largely structured production around a capital-intensive, industrial model. The declining price of computation, however, has inverted the capital structure of information and cultural production. Inexpensive desktop PCs, as well as digital video and audio systems, are now capable of performing most of the physical capital functions that once required substantial investments. Where physical capital costs for fixation and communication are low and widely distributed, and where existing information is itself a public good, the primary remaining scarce resource is human creativity. And it is under these conditions that the relative advantages of peer production emerge to much greater glory than possible before.

This leaves the motivation and organization questions. These generally would fall under the “tragedy of the commons” critique, which I purposefully invoke by calling the phenomenon “commons-based” peer production. The traditional objections to the commons are primarily twofold. First, no one will invest in a project if they cannot appropriate its benefits. That is, motivation will lack. Second, no one has the power to organize collaboration in the use of the resource. That is, organization will lack and collaboration will fail. The past decade or so, however, has seen an important emerging literature on some successful commons and common property regimes.<sup>17</sup> These primarily involve the introduction of a variety of nonproperty-based schemes for structuring cooperation among relatively limited groups of participants. While offering important insights into how formal and informal norms can structure collaboration, these studies of

---

16. While the reference to information as a public good is common, the reference to culture is not. I have no intention to go into subtle definitions of culture here, though I tend to follow the approach offered in J.M. BALKIN, CULTURAL SOFTWARE (1998), by thinking of culture as a framework for comprehension. By “culture” I mean a set of representations, conceptions, interpretations, knowledge of social behavior patterns, etc., whose particular application to reducing uncertainty for human action is too remote to be called “information,” but which is indispensable to the way we make sense of the world. “Cultural production” as I use it here can be done by parents, teachers, Hollywood, Mozart, the Pope, peer groups, and the guys playing guitars in Washington Square Park. Defined as a set of conceptions and their representations and as sets of behavioral instructions, its economic character is similar to ideas or information. Obviously, embodiments of culture, like a specific statue or building, are no more nonrival than embodiments of any other form of information, like a book or a corkscrew.

17. For discussions of commons, see ELINOR OSTROM, GOVERNING THE COMMONS: THE EVOLUTION OF INSTITUTIONS FOR COLLECTIVE ACTION (1990); and Carol Rose, *The Comedy of the Commons: Custom, Commerce, and Inherently Public Property*, 53 U. CHI. L. REV. 711 (1986). A brief discussion of these concepts as applied to peer production follows below. See *infra* notes 112-120 and accompanying text.

2002]

Coase's Penguin

common appropriation regimes do not give a complete answer to the sustainability of motivation and organization for the truly open, large-scale nonproprietary peer production projects we see on the Internet.

My answer to these problems occupies Part III. The motivation problem is solved by two distinct analytic moves. The first involves the proposition that diverse motivations animate human beings, and, more importantly, that there exist ranges of human experience in which the presence of monetary rewards is inversely related to the presence of other, social-psychological rewards. The interaction between money, love, and sex offers an obvious and stark example, but the tradeoffs that academics face between selling consulting services, on the one hand, and writing within a research agenda respected by peers, on the other hand, are also reasonably intuitive. Given these propositions, it becomes relatively straightforward to see that there will be conditions under which a project that can organize itself to offer social-psychological rewards removed from monetary rewards will attract certain people, or at least certain chunks of people's days, that monetary rewards would not.

The second analytic move involves understanding that when a project of any size is broken up into little pieces, each of which can be performed by an individual in a short amount of time, the motivation to get any given individual to contribute need only be very small. This suggests that peer production will thrive where projects have three characteristics. First, they must be modular. That is, they must be divisible into components, or modules, each of which can be produced independently of the production of the others. This enables production to be incremental and asynchronous, pooling the efforts of different people, with different capabilities, who are available at different times. Second, the granularity of the modules is important and refers to the sizes of the project's modules. For a peer production process to pool successfully a relatively large number of contributors, the modules should be predominately fine-grained, or small in size. This allows the project to capture contributions from large numbers of contributors whose motivation levels will not sustain anything more than small efforts toward the project. Novels, for example, at least those that look like our current conception of a novel, are likely to prove resistant to peer production.<sup>18</sup> In addition, a project will likely be more efficient if it can accommodate variously sized contributions. Heterogeneous granularity will

---

18. The most successful novel-like enterprise on the Internet of which I know is "The Company Therapist." Pipsqueak Prods., *The Company Therapist*, at <http://www.thetherapist.com> (last visited Aug. 31, 2002). There, the collaborative fiction problem was solved by building a system that enabled anyone to contribute a small chunk—patient's interview notes, therapist's comments, etc.—to the company therapist's files. The common project is to create a fascinating mosaic of people and stories seen through the eyes of a company therapist. Most collaborative fiction sites, however, suffer from the fact that modularity and granularity lead to disjunction relative to our expectations of novels.

allow people with different levels of motivation to collaborate by making smaller- or larger-grained contributions, consistent with their levels of motivation. Third, and finally, a successful peer production enterprise must have low-cost integration, which includes both quality control over the modules and a mechanism for integrating the contributions into the finished product. If a project cannot defend itself from incompetent or malicious contributions and integrate the competent modules into a finished product at sufficiently low cost, integration will either fail or the integrator will be forced to appropriate the residual value of the common project—usually leading to a dissipation of the motivations to contribute *ex ante*. Automated integration and iterative peer production of integration are the primary mechanisms by which peer production projects described in this Article have lowered the cost of integration to the point where they can succeed and sustain themselves. The use of free software to integrate peer production of other information goods is a prime example. As for a project’s mechanisms for defending itself from incompetent or malicious contributions, one sees peer production enterprises using a variety of approaches toward solving collective action problems that are relatively familiar from the commons literature offline. These include various formal rules, like the GNU General Public License (GPL)<sup>19</sup> that prevents defection<sup>20</sup> from many free software projects, including most prominently its flagship, GNU/Linux. They also include technical constraints that prevent or limit the effects of defection. Social norms too play a role in sustaining some of these collaborations, in both small groups and larger groups where the platform allows for effective monitoring and repair when individuals defect. Finally, the sheer size of some of these projects enables the collaboration platform to correct for defection by using redundancy of

---

19. The GNU GPL is the most important institutional innovation of the Free Software Foundation founded by Richard Stallman. Free Software Found., *GNU General Public License*, at <http://www.fsf.org/copyleft/gpl.html> (last visited Aug. 31, 2002). It prevents defection from free software projects in the form of combining code others have written with one’s own code and then releasing it under more restrictive license terms than the original free software. This license does not prevent commercial distribution of free software for a fee. It places certain limits on how the software can be used as an input into derivative works that would be made less free than the original. In this, it radically breaks from the concept of the public domain that underlies copyright law’s general background rule for nonproprietary materials. For discussions of the GPL and its legal nature and institutional characteristics, see Eben Moglen, *Enforcing the GPL*, LINUXUSER, Sept. 2001, at 66, [http://www.linuxuser.co.uk/articles/issue14/lu14-Free\\_Speech-Enforcing\\_the\\_GPL.pdf](http://www.linuxuser.co.uk/articles/issue14/lu14-Free_Speech-Enforcing_the_GPL.pdf); and Eben Moglen, *Enforcing the GPL II*, LINUXUSER, Oct. 2001, at 66, [http://www.linuxuser.co.uk/articles/issue15/lu15-Free\\_Speech-Enforcing\\_the\\_GPL\\_part\\_two.pdf](http://www.linuxuser.co.uk/articles/issue15/lu15-Free_Speech-Enforcing_the_GPL_part_two.pdf). Moglen’s views are particularly important since he has been General Counsel to the Free Software Foundation for the past decade and has more experience with enforcing this license than anyone else. More detailed academic treatments include McGowan, *supra* note 13, and Margaret Jane Radin & R. Polk Wagner, *The Myth of Private Ordering: Rediscovering Legal Realism in Cyberspace*, 73 CHI.-KENT L. REV. 1295 (1998).

20. I use the term “defection” to describe any action that an agent who participates in a cooperative enterprise can take to increase his or her own benefit from the common effort in a way that undermines the success or integrity of the common effort.

2002]

Coase's Penguin

contributions and averaging out of outliers—be they defectors or incompetents.

The normative implications of recognizing peer production are substantial. At the level of political morality, the shape of freedom and equality in the emerging social-technological condition we associate with the Internet is at stake. Political views can take radical forms, both anarchistic and libertarian, as they do in the work of Eben Moglen, who was the first to identify the phenomenon I now call peer production,<sup>21</sup> and in the minds of many in the free software community.<sup>22</sup> But the stakes for freedom and equality are high for a wide range of liberal commitments.<sup>23</sup> At the level of institutional design, the emergence of commons-based peer production adds a new and deep challenge to the prevailing policy of rapid expansion of the scope of exclusive rights in information and culture that has been the predominant approach in the past twenty-five years, as James Boyle's work on the second enclosure movement elegantly elucidates.<sup>24</sup> Additionally, the dynamic of decentralized innovation plays a central role in Lawrence Lessig's forceful argument for embedding the openness of commons in the architecture of the Internet.<sup>25</sup> In this Article, however, I do not attempt to add to the normative literature. Instead, the Article is intended as a purely descriptive account of the scope of the empirical phenomenon and its analytic drivers.

One important caveat is necessary. I am *not* suggesting that peer production will supplant markets or firms. I am *not* suggesting that it is always the more efficient model of production for information and culture. What I am saying is that this emerging third model is (1) distinct from the other two and (2) has certain systematic advantages over the other two in identifying and allocating human capital/creativity. When peer production

---

21. Eben Moglen, *Anarchism Triumphant: Free Software and the Death of Copyright*, 4 FIRST MONDAY 1 (Aug. 2, 1999), at [http://www.firstmonday.dk/issues/issue4\\_8/moglen/](http://www.firstmonday.dk/issues/issue4_8/moglen/). The descriptive insight in that paper that corresponds to peer production is the phenomenon he calls Moglen's Metaphorical Corollary to Faraday's Law:

Moglen's Metaphorical Corollary to Faraday's Law says that if you wrap the Internet around every person on the planet and spin the planet, software flows in the network. It's an emergent property of connected human minds that they create things for one another's pleasure and to conquer their uneasy sense of being too alone.

*Id.*

22. Canonical, of course, are Richard Stallman's ideas, which permeate the "Philosophy of the GNU Project." See Free Software Found., Philosophy of the GNU Project, at <http://www.gnu.org/philosophy/philosophy.html> (last visited Aug. 25, 2002).

23. I outline the breadth of the range of liberal commitments affected by these issues in Yochai Benkler, *Freedom in the Commons: Towards a Political Economy of Information*, 52 DUKE L.J. (forthcoming 2003).

24. See James Boyle, The Second Enclosure Movement and the Construction of the Public Domain, Paper Presented at the Conference on the Public Domain, Duke Law School, Nov. 9-11, 2001, at <http://www.law.duke.edu/pd/papers/boyle.pdf>.

25. LAWRENCE LESSIG, THE FUTURE OF IDEAS: THE FATE OF THE COMMONS IN A CONNECTED WORLD (2001).

will surpass the advantages that the other two models may have in triggering or directing human behavior through the relatively reliable and reasonably well-understood triggers of money and hierarchy is a matter for more detailed study. I offer some ways of understanding the limitations of this model of production in Part III, but do not attempt a full answer to these questions here.

### I. PEER PRODUCTION ALL AROUND

While open source software development has captured the attention and devotion of many, it is by no stretch of the imagination the first or most important instance of production by peers who interact and collaborate without being organized on either a market-based or a managerial/hierarchical model. Most important in this regard is the academic enterprise, and in particular scientific research. Thousands of individuals make contributions to a body of knowledge, set up internal systems of quality control, and produce the core of our information and knowledge environment. These individuals do not expect to exclude from their product anyone who does not pay for it, and for many of them the opportunity cost of participating in academic research, rather than applying themselves to commercial enterprise, carries a high economic price tag. In other words, individuals produce on a nonproprietary basis and contribute their product to a knowledge “commons” that no one is understood as “owning,” and that anyone can, indeed is required by professional norms to, take and extend. We appropriate the value of our contributions using a variety of methods: service-based rather than product-based models (teaching rather than book royalties), grant funding from government and nonprofit sources, as well as reputation and similar intangible but immensely powerful motivations embodied in prizes, titles, etc. In the excitement of a moment that feels like one of great transformation, it is easy, though unjustifiable, to forget that information production is one area where we have always had a mixed system of commercial/proprietary and nonproprietary peer production—not as a second best or a contingent remainder from the Middle Ages, but because at some things the nonproprietary peer production system of the academic world is simply better.<sup>26</sup>

---

26. An early version of this position is Richard R. Nelson, *The Simple Economics of Basic Scientific Research*, 67 J. POL. ECON. 297 (1959). More recently one sees the work, for example, of Rebecca S. Eisenberg. See Rebecca S. Eisenberg, *Public Research and Private Development: Patents and Technology Transfer in Government-Sponsored Research*, 82 VA. L. REV. 1663, 1715-24 (1996) (summarizing the argument for giving preference to universities and nonprofit institutions in the allocation of patent rights). For a historical description of the role of market and nonmarket institutions in science, see P.A. David, *From Market Magic to Calypso Science Policy*,

2002]

Coase's Penguin

In one way, however, academic peer production and commercial production are similar. Both are composed of people who are professional information producers. The individuals involved in production have to keep body and soul together from information production. However low the academic salary is, it must still be enough to permit one to devote most of one's energies to academic work. The differences reside in the modes of appropriation and in the modes of organization—in particular, how projects are identified and how individual effort is allocated to projects. Academics select their own projects and contribute their work to a common pool that eventually comprises our knowledge of a subject matter, while nonacademic producers will often be given their marching orders by managers, who take their cue from market studies and eventually sell the product in the market.

Alongside the professional model, it is also important to recognize that we have always had nonprofessional information and cultural production on a nonproprietary model. Individuals talking to each other are creating information goods, sometimes in the form of what we might call entertainment and sometimes as a means for news distribution or commentary. Nonprofessional production has been immensely important in terms of each individual's information environment. If one considers how much of the universe of communications one receives in a day from other individuals in one-to-one or small-scale interactions—such as e-mail, lunch, or hallway conversations—the effect becomes tangible.

Nonetheless, ubiquitous computer communications networks are bringing about a dramatic change in the scope, scale, and efficacy of peer production. As computers and network connections become faster, cheaper, and more ubiquitous, we are seeing the phenomenon of nonprofessional peer production of information scale to much larger sizes and perform more complex tasks than were possible in the past for nonprofessional production. To make this phenomenon more tangible, I will describe in this Part a number of such enterprises, organized to demonstrate the feasibility of this approach throughout the information-production and information-exchange chain.

While it is possible to break an act of communication into finer-grained subcomponents,<sup>27</sup> largely we see three distinct functions involved in the process. First, there is an initial utterance of a humanly meaningful statement. Writing an article or drawing a picture is such an action, whether done by a professional or an amateur and whether high quality or low. Second, there is the separate function of mapping the initial utterances onto

---

26 RES. POL'Y 229 (1997) (reviewing TERENCE KEALEY, THE ECONOMIC LAWS OF SCIENTIFIC RESEARCH (1996)).

27. See Yochai Benkler, *Communications Infrastructure Regulation and the Distribution of Control over Content*, 22 TELECOMM. POL'Y 183, 186 (1998).

a knowledge map. In particular, an utterance must be understood as “relevant” and “credible.” Relevance is a subjective question of mapping an utterance on the conceptual map of a given user seeking information for a particular purpose defined by that individual. If I am interested in learning about the political situation in Macedonia, a news report from Macedonia or Albania is relevant, even if sloppy, while a Disney cartoon is not, even if highly professionally rendered. Credibility is a question of quality by some objective measure that the individual adopts for purposes of evaluating a given utterance. Again, the news report may be sloppy and not credible, while the Disney cartoon may be highly accredited *as a cartoon*. The distinction between the two is somewhat artificial, however, because very often the utility of a piece of information will depend on a combined valuation of its credibility and relevance. A *New York Times* story on the Balkans in general, for example, will likely be preferable to excited gossip in the cafeteria specifically about Macedonia. I will therefore refer to “relevance/accreditation” as a single function for purposes of this discussion, keeping in mind that the two are complementary and not entirely separable functions that an individual requires in order to use utterances of others in putting together the user’s understanding of the world.

Finally, there is the function of distribution, or how one takes an utterance produced by one person and distributes it to other people who find it credible and relevant. In the mass media world, these functions were often, though not always, integrated. NBC News produced the utterances, gave them credibility, and distributed them. The Internet is permitting much greater disaggregation of these functions, and so this Part will proceed to describe how each component of this information-production chain is being produced using a peer-based model on the Internet for information and cultural goods other than software.

#### A. Content

NASA Clickworkers is “an experiment that showed that public volunteers (clickworkers), many working for a few minutes here and there . . . can do some routine science analysis that would normally be done by a scientist or graduate student working for months on end.”<sup>28</sup> Users can mark craters on maps of Mars, classify craters that have already been marked, or search the landscape of Mars for “honeycomb” terrain. The project is “a pilot study with very limited funding, run part-time by one

---

28. NASA, Welcome to the Clickworkers Study, at <http://clickworkers.arc.nasa.gov/top> (last visited Aug. 31, 2002).

2002]

Coase's Penguin

software engineer, with occasional input from two scientists.<sup>29</sup> In its first six months of operation, more than 85,000 users visited the site, with many contributing to the 1.9 million entries recorded (including redundant entries of the same craters used to average out errors). An analysis of the quality of markings showed that "the automatically-computed consensus of a large number of clickworkers is virtually indistinguishable from the inputs of a geologist with years of experience in identifying Mars craters."<sup>30</sup> The tasks performed by clickworkers are discrete, and each iteration is easily performed in a matter of minutes. As a result, users can choose to work for a few minutes by doing one iteration or for hours by doing many, with an early study of the project suggesting that some clickworkers indeed work on the project for weeks, but that 37% of the work was done by one-time contributors.<sup>31</sup>

The Clickworkers project is a particularly crisp example of how complex professional tasks that formerly required budgeting the full-time salaries of a number of highly trained individuals can be reorganized to be performed by tens of thousands of volunteers in increments so minute that the tasks can now be performed on a much lower budget. This low budget is devoted to coordinating the volunteer effort, and the raw human capital needed is contributed for the fun of it. The professionalism of the original scientists is replaced by a combination of fine-grained modularization of the task with redundancy and automated averaging out of both errors and purposeful defections (for example, purposefully erroneous markings).<sup>32</sup> NASA scientists running this experiment tapped into a vast pool of five-minute increments of human judgment applied with motivation that is unrelated to keeping together the bodies and souls of the agents.

While Clickworkers is a distinct, self-conscious experiment, it suggests characteristics of distributed production that are quite widely observable. Consider, for example, how the networked environment has enabled new ways of fulfilling the traditional function of encyclopedias or almanacs. At the most general level, consider the World Wide Web itself. Individuals put up websites with all manner of information, in all kinds of quality and focus, for reasons that have nothing to do with external, well-defined economic motives—just like the individuals who identify craters on Mars. A user interested in information need only plug a request into a search

---

29. NASA, Credits and Contacts, at <http://clickworkers.arc.nasa.gov/contact> (last visited Aug. 31, 2002).

30. NASA, CLICKWORKERS RESULTS: CRATER MARKING ACTIVITY (July 3, 2001), at <http://clickworkers.arc.nasa.gov/documents/crater-marking.pdf>.

31. B. KANEFSKY ET AL., CAN DISTRIBUTED VOLUNTEERS ACCOMPLISH MASSIVE DATA ANALYSIS TASKS? 1 (2001), at <http://clickworkers.arc.nasa.gov/documents/abstract.pdf>.

32. NASA, *supra* note 30, at 3 (describing, among other things, the exclusion of the markings of a student in an art class who marked concentric circles for a class assignment instead of trying to mark craters).

engine like Google, and dozens or hundreds of websites will appear. Now, there is a question of how to select among them—the question of relevance and accreditation—but that is for the next Section. For now it is important to recognize that the Web is a global library produced by millions of people. Whenever you sit down to search for information, there is a very high likelihood that someone, somewhere, has produced a usable answer, for whatever reason—pleasure, self-advertising, or fulfilling some other public or private goal as a nonprofit or for-profit institution that sustains itself by means other than selling the information you need. The power of the Web to answer such an encyclopedic question comes not from the fact that one particular site has all the great answers. It is not an Encyclopedia Britannica. The power comes from the fact that it allows a user looking for specific information at a given time to collect answers from a sufficiently large number of contributions. The task of sifting and accrediting falls to the user, who is motivated by the need to find an answer to the question posed. As long as there are tools to lower the cost of *that* task to a level acceptable to the user, the Web will have “produced” the information content the user sought. These are not trivial considerations, but they are also not intractable. As we shall see, some of the solutions can themselves be peer produced and some solutions are emerging as a function of the speed of computation and communication, which enables more efficient technological solutions.

One might argue that the Internet is still not an encyclopedia, in the sense of a coherently ordered locus of a wide range of human knowledge in relatively accessible and digested form. Can that task, which requires more disciplined writing, be performed within a distributed model? The beginning of an answer is provided by the Wikipedia project.<sup>33</sup> The project involves roughly 2000 volunteers who are collaborating to write an encyclopedia. The project runs on a free software collaborative authorship tool, Wiki, which is a markup language similar in concept to HTML, but is relatively easier to implement, allows multiple people to edit a single document and interlock multiple documents, and generates archives of the changes made to each. While 2000 people have not been able to generate a complete encyclopedia in roughly 18 months of operation, they have made substantial progress, producing about 30,000 articles, and readers are invited to evaluate the quality.<sup>34</sup> A comparison to www.encyclopedia.com, the online version of the Columbia Encyclopedia, would suggest that Wikipedia cannot yet be said to be either systematically better or worse. Given that it is a volunteer effort, and that the comparison is to an

---

33. Wikipedia, Main Page, at <http://www.wikipedia.com> (last visited Aug. 31, 2002).

34. The terms “chimpanzee,” “computational complexity theory,” or simply “copyright,” for example, provide good demonstrations.

2002]

Coase's Penguin

established commercial encyclopedia, that is actually saying quite a bit. Perhaps the most interesting characteristic about Wikipedia is the self-conscious social-norms-based dedication to objective writing. The following fragments from the self-described essential characteristics and basic policies of Wikipedia are illustrative:

First and foremost, the Wikipedia project is self-consciously an encyclopedia—rather than a dictionary, discussion forum, web portal, etc. See encyclopedia as well as what Wikipedia is not.

....

Wikipedia's participants commonly follow, and enforce, a few basic policies that seem essential to keeping the project running smoothly and productively. The following are just a few of those policies; for more information, please see Wikipedia policy.

First, because we have a huge variety of participants of all ideologies, and from around the world, Wikipedia is committed to making its articles as unbiased as possible. The aim is not to write articles from a single *objective* point of view—this is a common misunderstanding of the policy—but rather, to fairly and sympathetically present all views on an issue. See neutral point of view page for further explanation, and for a very lengthy discussion.<sup>35</sup>

The participants of Wikipedia are plainly people who like to write. Some of them participate in other collaborative projects, like Everything2.com.<sup>36</sup> But when they enter the common project of Wikipedia,

---

35. Wikipedia, Essential Characteristics, at <http://www.wikipedia.com/wiki/wikipedia> (last visited Aug. 31, 2002) (hyperlinks indicated by underlining). The “Neutral Point of View” page is indeed revealing of how explicit and central to the project the social norm of objective writing is. See Wikipedia, Neutral Point of View, at <http://www.wikipedia.com/wiki/wikipedia:neutral+point+of+view> (last visited Aug. 31, 2002).

36. See Everything Dev. Co., Everything@Everything2.com, at <http://www.everything2.com> (last visited Aug. 31, 2002). Everything2 is a “complex online community with a focus to write, publish and edit a quality database of information, insight and humor.” *Id.* (under the “Everything FAQ” hyperlink). The system enables registered users to post “write-ups” and create “nodes” pertaining to particular topics that they define. It does not have a directory structure; instead, nodes are linked together with hypertext within the text of the node and also with a matrix of related links at the bottom of each node. The linking is done initially by the author—thereby self-generating a conceptual map—and later by others. A node is a particular topic identified by the title of the node. After the author of the first write-up creates a “nodedshell,” other users can add additional write-ups to that node. Write-ups are constantly being reviewed and removed by editors. Editors are chosen based on “merit, seniority and writing skill.” Everything Dev. Co., The Power Structure of Everything2, at [http://www.everything2.com/index.pl?node\\_id=743129](http://www.everything2.com/index.pl?node_id=743129) (last visited Nov. 11, 2002). Everything2 also contains a voting system for non-editor users to vote on each other’s write-ups. Although each write-up has a reputation based on whether it has been voted up or down, the write-up does not get automatically filtered due to a low reputation. In other

they participate in a particular way—a way that the group has adopted to make its “encyclopedia.”<sup>37</sup> Wikipedia provides a rich example of a medium-sized collection of individuals who successfully collaborate to create an information product of mid- to highbrow quality. In particular, it suggests that even in a group of this size, social norms coupled with a simple facility to allow any participant to edit out blatant opinion written by another in contravention of the social norms keep the group on track.

Perhaps the most sophisticated locus of peer reviewed, mid- to high-quality essays published on the Internet as of early 2002 is Kuro5hin, also known as K5.<sup>38</sup>

Kuro5hin.org is a community of people who like to think. You will not find garbage in the discussions here, because noise is not tolerated. This is a site for people who want to discuss the world they live in. It’s a site for people who are on the ground in the modern world, and who sometimes look around and wonder what they have wrought.<sup>39</sup>

As of March 2002, it appeared that Kuro5hin had roughly 25,000 users.<sup>40</sup> Articles run a broad gamut of topics but are supposed to be roughly centered around technology and culture. The general headings include Technology, Culture, Politics, Media, News, Op-ed, Columns, Meta (dedicated to discussion of K5 itself), and MLP (mindless link propagation, a general catchall category of things the community members find interesting). The articles include news reports from other sources, but most of the interesting materials provide some form of commentary as well. The articles and responses to them are fairly substantial.

---

words, the system combines individually authored materials and individually defined mappings of relevance of materials with common procedures, some purely democratic and some based on a rotating hierarchy of editors appointed by experience and reputation built from the collective judgments of their peers. The result is a substantial database of writings on a wide variety of topics.

37. On their interpretation, creating an encyclopedia entry means conveying in brief terms the state of the art on the topic, including divergent opinions about it, but not the author’s opinion. Whether that is an attainable goal is a subject of interpretive theory and is a question as applicable to a professional encyclopedia as it is to Wikipedia.

38. Kuro5hin, Front Page, at <http://www.kuro5hin.org> (last visited Aug. 31, 2002). The discussion here is deeply indebted to the work of Caio M.S. Pereira Neto. See Caio M.S. Pereira Neto, Kuro5hin.org, Collaborative Media, and Political Economy of Information (May 24, 2002) (unpublished manuscript, on file with author). Another source is Everett Teach et al., Ethnography of Kuro5hin.org, at <http://ccwf.cc.utexas.edu/~hackett/k5/> (last visited Aug. 31, 2002).

39. Kuro5hin, Mission Statement, at <http://www.kuro5hin.org/special/mission> (last visited Sept. 25, 2002) (emphasis omitted).

40. How Will K5 Avoid Being Crushed by Content? (Mar. 17, 2002), at <http://www.kuro5hin.org/story/2002/3/16/51221/8976>.

2002]

## Coase's Penguin

The site and community have a heavy emphasis on the quality of materials published. The guide to article submissions<sup>41</sup> emphasizes quality of information and writing multiple drafts and prepares new contributors for the experience of close peer review of their submission. Additionally, the software that runs Kuro5hin, Scoop, a free software project initiated by one of the cofounders of K5, implements a series of steps both before and after submission and publication of an article that serve as collaborative quality-control mechanisms. The emphasis on quality is enforced by the site's mechanism for peer review prepublication and peer commentary postpublication.<sup>42</sup>

A very different type of trend in collaborative creation is the emergence and rise of computer games, in particular multiplayer online games. These fall in the same cultural "time slot" as television shows and movies of the twentieth century. The interesting thing about them is that they are structurally different. In a game like Ultima Online or EverQuest, the role of the commercial provider is not to tell a finished, highly polished story to be consumed start to finish by passive consumers. Rather, the role of the game provider is to build tools with which users collaborate to tell a story. There have been observations about this approach for years regarding MUDs and MOOs.<sup>43</sup> The point here is that there is a type of "content" that

---

41. Kuro5hin, FAQ—Article Submission Questions, at <http://www.kuro5hin.org/?op=special;page=article> (last visited Aug. 31, 2002).

42. When an article is submitted, it is not automatically placed in a publicly viewable space. It is placed, instead, in a submission queue. At that point, all registered users of K5 have an opportunity to comment on the article, provide suggestions for correction and improvement, and vote their opinion whether they think the story should be placed on the front page, a specialty page, or rejected. The system determines some critical number of votes necessary for any one of these actions, based on the number of users then registered. Typically, rejection requires fewer votes than acceptance. Articles may be resubmitted after being rejected, typically after having been revised in accordance with the comments. The system up to this point is remarkably similar to academic peer review in many respects, except for the scope of participation and the egalitarian and democratic structure of the editorial decision. After publication, K5 provides the platform for readers to comment on articles and for other readers to rate these comments for their relevance and quality. The system is different in various respects from the Slashdot system described in detail *infra* Section I.B, but the principle is the same. It permits readers to post comments, and it permits other readers to rate comments as better or worse. It aggregates these individual ratings into collective judgments about the quality of comments, judgments that can then be used by the site's readers to filter out lower-quality comments. In general, all these characteristics go to questions of how one generates relevance and accreditation on a peer production model and will be explored in greater detail in the context of other sites in the next Section. The point to take away at this juncture is that part of what makes K5 so successful in maintaining quality is a rather elaborate, large-scale peer review system and postpublication commentary, which itself is then peer reviewed in an iterative process.

43. MUDs (Multi-User Dungeon or Multi-User Dimension) and MOOs (MUD, Object Oriented) are acronyms for software programs that create an interactive multiuser networked text-based virtual world. The software maintains a database of users and objects with which the users can interact in a variety of ways. MUDs are typically built around a theme. MUD "worlds" are often based on books, movies, cartoons, and other role-playing games. See RPer's Resources for Interactive Roleplaying on Moos/Mucks/Muds/Mushes, at <http://www.geocities.com/TimesSquare/9944/> (last visited Aug. 31, 2002). Pavel Curtis, creator of perhaps the most famous

---

of MOOs, LambdaMOO, identified three elements that distinguish MUDs from typical role-playing games:

- A MUD is not goal-oriented; it has no beginning or end, no “score,” and no notion of “winning” or “success.” In short, even though users of MUDs are commonly called *players*, a MUD isn’t really a game at all.
- A MUD is extensible from within; a user can add new objects to the database such as rooms, exits, “things,” and notes. Certain MUDs, including the one I run, even support an embedded programming language in which a user can describe whole new kinds of behavior for the objects they create.
- A MUD generally has more than one user connected at a time. All of the connected users are browsing and manipulating the same database and can encounter the new objects created by others. The multiple users on a MUD can communicate with each other in real time.

Pavel Curtis, *Mudding: Social Phenomena in Text-Based Virtual Realities*, in PROCEEDINGS OF THE 1992 CONFERENCE ON THE DIRECTIONS AND IMPLICATIONS OF ADVANCED COMPUTING (2002), at <http://citeseer.nj.nec.com/curtis92mudding.html>. There are acronyms for MUD-like variations including MUSH, MUX, and MUCK. All of the variations run basically the same software; the primary difference between them is how much freedom the characters have to modify the environment. All M\*s (M\* refers to any MUD-like variants) are administered in some way by those who set up the software and maintain the connectivity. Typically, the administrator will set up the initial world and implement some coded commands. The administrator will also set up a hierarchy of user levels granting users more control over the objects within the game as they advance, such as the ability to create coded commands. It is these decisions—how much of the world does the administrator create, how rich are the coded commands, does the administrator allow users to have the power to manipulate the game—that distinguish the various M\*s from each other. MUDs are typically heavy on coded commands and designed to be battle-ready. MUSHes, on the other hand, are “unlikely to have coded commands to the same extent that a MUD will, relying instead on arbitration or consent to determine the effects of actions.” Michael Sullivan, An Explanation of Terminology, at <http://wso.williams.edu:8000/~msullivan/mushes/explan.html> (last visited Aug. 31, 2002). MOOs are perhaps the exception in that most of them are not role-playing, but “educational or social.” *Id.*

Most important in the history of MUDs was LambdaMOO. “LambdaMOO is a MOO: a MUD that uses an object-oriented programming language to manipulate objects in the virtual world.” AT&T Cobot Project, What’s a LambdaMOO?, at <http://cobot.research.att.com/lambdamoo.html> (last visited Aug. 31, 2002). LambdaMOO was created in 1990 by Pavel Curtis as a social experiment. “[It] is the first, most diverse, oldest, largest, and most well-known MOO.” Rebecca Spainhower, Virtually Inevitable: Real Problems in Virtual Communities, at <http://world.std.com/~rs/inevitable.html> (last visited Aug. 31, 2002). “When Pavel Curtis took on the project of developing the MOO environment, he gave it a social focus instead of the game goal of traditional MUDs.” *Id.* The original site has remained active for over a decade and continues to thrive with over 100,000 people having participated in this one virtual world. As a result, LambdaMOO is

a long-standing, ongoing experiment in collective programming and creation, with often stunning results that can only be fully appreciated firsthand. Inventions include technical objects, such as the lag meter, which provides recent statistics on server load; objects serving a mix of practical and metaphorical purposes, such as elevators that move users between floors; objects with social uses, such as the birthday meter, where users register their birthdays publicly; and objects that just entertain or annoy, such as the Cockatoo, a virtual bird who occasionally repeats an utterance recently overheard.

AT&T Cobot Project, *supra* (emphasis omitted). Jennifer Mnookin and Rebecca Spainhower detail the evolution of the social structure of LambdaMOO. See Jennifer L. Mnookin, *Virtual(l)y) Law: The Emergence of Law in LambdaMOO*, 2 J. COMPUTER-MEDIATED COMM. 8 (June 1996), at <http://www.ascusc.org/jcmc/vol2/issue1/lambdamoo.html>; Spainhower, *supra*. Generally, the MOO was administered by a few system administrators (called wizards within the game). Haakon (Pavel Curtis’s wizard character) drafted a set of guidelines for behavior. When administration became too overwhelming for the wizards, they appointed an “Architecture Review Board” of fifteen trusted users to allocate space to new users. The wizards were still responsible, however, for

2002]

Coase's Penguin

can be produced in a centralized professional manner—the screenwriter here replaces the scientist in the NASA Clickworkers example—that can also be organized using the appropriate software platform to allow the story to be written by the many users as they experience it. The users are coauthors whose individual contributions to the storyline are literally done for fun. They are playing a game, but they are spending real economic goods—their attention and substantial subscription fees—on a form of entertainment that displaces passive reception of a finished, professionally manufactured good with active coproduction of a storyline. The individual contributions are much more substantial than the time needed to mark craters, but then the contributors are having a whole lot more fun manipulating the intrigues of their imaginary Guild than poring over digitized images of faint craters on Mars.

#### B. *Relevance/Accreditation*

You might say that many distributed individuals can produce content, but that it is gobbledegook. Who in their right mind wants to get answers to legal questions from a fifteen-year-old child who learned the answers from watching Court TV?<sup>44</sup> The question then becomes whether relevance and accreditation of initial utterances of information can itself be produced on a peer production model. The answer is that it can. Some of the most prominent web-based enterprises, both commercial and noncommercial, demonstrate this answer by breaking off precisely the accreditation and relevance piece of their product for peer production. Amazon.com and Google are good examples in the commercial arena.

Amazon uses a mix of mechanisms to highlight books and other products that its users are likely to buy.<sup>45</sup> A number of these mechanisms

---

dealing with unruly users and mediating disputes. In 1993, the wizards turned that responsibility over to the community at large by implementing a democratic petitioning and balloting system. Since that time, the community has addressed problems of population growth, harassment, and the behavior of anonymous guest accounts.

44. Michael Lewis, *Faking It*, N.Y. TIMES, July 15, 2001, § 6 (Magazine), at 32.

45. See Amazon.com, at <http://www.amazon.com> (last visited Aug. 31, 2002) (under the “Friends and Favorites” hyperlink). Amazon is constantly testing new methods of peer producing relevance and accreditation mechanisms and removing unpopular methods. These include both automatically generated and human-made relevance maps. For example, “Page You Made” is based on the user’s recent clicks on the site and lists a “Featured Item” as well as several “Quick Picks,” which are products that are similar to the recently viewed items. The page features “Listmania” lists, which are user-created topical lists, and a “More To Explore” section that provides relevant links to a topical directory of the Amazon inventory. Users can also “Share Purchases” and make their purchases available for other users to see. If the user finds a person with similar tastes, these options could aid with relevance, and if the user finds a particularly trustworthy person, it could aid in accreditation of the product. Amazon also provides discussion boards for direct exchange between users. Amazon creates “Purchase Circles,” which are “highly specialized bestseller lists,” based on aggregated data divided either geographically (by town or city) or organizationally (by schools, government offices, or corporations). *Id.* (under the “Friends

produce relevance and accreditation by harnessing the users themselves. At the simplest level, the recommendation “Customers Who Bought Items You Recently Viewed Also Bought These Items” is a mechanical means of extracting judgments of relevance and accreditation from the collective actions of many individuals who produce the datum of relevance as a by-product of making their own purchasing decisions. At a more self-conscious level (self-conscious, that is, on the part of the user), Amazon allows users to create topical lists, and to track other users as their “Friends and Favorites,” whose decisions they have learned to trust. Amazon also provides users with the ability to rate books they buy, generating a peer-produced rating by averaging the ratings. The point to take home from Amazon is that a corporation that has done immensely well at acquiring and retaining customers harnesses peer production to provide one of its salient values—its ability to allow users to find things they want quickly and efficiently.

Similarly, Google, which is widely recognized as the most efficient general search engine currently operating, introduced a crucial innovation into ranking results that made it substantially better than any of its competitors. While Google uses a text-based algorithm to retrieve a given universe of web pages initially, its PageRank software employs peer production of ranking in the following way.<sup>46</sup> The engine treats links from other websites pointing to a given website as votes of confidence. Whenever one person’s page links to another page, that person has stated quite explicitly that the linked page is worth a visit. Google’s search engine counts these links as votes of confidence in the quality of that page as compared to other pages that fit the basic search algorithm. Pages that themselves are heavily linked-to count as more important votes of confidence, so if a highly linked-to site links to a given page, that vote counts for more than if an obscure site links to it. By doing this, Google harnessed the distributed judgments of many users, with each judgment created as a by-product of making his or her own site useful, to produce a highly valuable relevance and accreditation algorithm. Google’s experience is particularly salient when juxtaposed with that of GoTo.com, which was a search engine that sold placement on the search result list to the highest

---

and Favorites” followed by the “Purchase Circles” hyperlinks). The data is analyzed and compared to site-wide trends to come up with lists of items that are more popular with that particular group than with the general population. If users find a list particularly useful, they can bookmark the list to view the changes as the list is updated to reflect new sales data. Amazon software also recommends certain products to the user. These “Recommendations” are based on items the user has purchased or rated, as well as their activity on the site contrasted with other users’ activity. As a result, the recommendations can change when the user purchases or reviews an item, or when the interests of other consumers change.

46. See Google, Inc., Our Search: Google Technology, at <http://www.google.com/technology/> (last visited Aug. 31, 2002).

2002]

## Coase's Penguin

bidder. It turns out that the site owner's willingness to pay *to be seen* is not necessarily a good measure of the utility its site provided to people who are searching the Web. Google recently replaced Overture, GoTo's current name, as America Online's (AOL's) default search engine.<sup>47</sup> A casual search using both will reveal the difference in quality between the two, and a search for "Barbie" will also yield interesting insights into the political morality of pricing as opposed to voting as the basis of relevance algorithms.

While Google is an automated mechanism of collecting human judgment as a by-product of some other activity, there are also important examples of distributed projects self-consciously devoted to peer production of relevance. Most prominent among these is the Open Directory Project (ODP).<sup>48</sup> The site relies on tens of thousands of volunteer editors to determine which links should be included in the directory. Acceptance as a volunteer requires application. Admission relies on a peer review process based substantially on seniority and the extent of a volunteer's engagement, as measured by the extent of his or her contributions. The site is hosted and administered by Netscape, which pays for server space and a small number of employees to administer the site and set up the initial guidelines, but licensing is free to the number of sites who use ODP as their web directory.<sup>49</sup> This presumably adds value to AOL and Netscape's commercial search engine and portal as well as enhances the company's goodwill. The volunteers are not affiliated with Netscape, receive no compensation, and manage the directory out of the joy of doing so or for other internal or external motivations. The volunteers spend time selecting sites for inclusion in the directory (in small increments of perhaps fifteen minutes per site reviewed), thereby producing a comprehensive, high-quality, human-edited directory of the Web—competing with, and quite possibly outperforming, Yahoo in this category.

Perhaps the most elaborate multilayer mechanism for peer production of relevance and accreditation is Slashdot.<sup>50</sup> Billed as "News for Nerds," Slashdot primarily consists of users commenting on initial submissions that cover a variety of technology-related topics. The submissions are typically a link to an off-site story, coupled with some initial commentary from the person who submits the piece. Users follow up the initial submission with comments that often number in the hundreds. The initial submissions and the approach to sifting through the comments of users for relevance and

---

47. David F. Gallagher, *AOL Shifts Key Contract to Google*, N.Y. TIMES, May 2, 2002, at C4.

48. Open Directory Project, *at* <http://www.dmoz.org> (last visited Aug. 31, 2002).

49. See Open Directory Project, Sites Using ODP Data, *at* [http://dmoz.org/Computers/Internet/Search/Directories/Open\\_Directory\\_Project/Sites\\_Using\\_ODP\\_Data/](http://dmoz.org/Computers/Internet/Search/Directories/Open_Directory_Project/Sites_Using_ODP_Data/) (last visited Sept. 24, 2002) (listing the sites that are currently using the ODP).

50. Open Source Dev. Network, Inc., *Slashdot: News for Nerds, Stuff That Matters*, *at* <http://slashdot.org> (last visited Aug. 31, 2002).

accreditation provide a rich example of how this function can be performed on a distributed, peer production model.

It is important initially to understand that the function of posting a story from another site onto Slashdot, the first “utterance” in a chain of comments on Slashdot, is itself an act of relevance production. The person submitting the story is telling the community of Slashdot users, “Here is a story that people interested in ‘News for Nerds’ should be interested in.” This initial submission of a link is itself filtered by “authors” (really editors), most of whom are paid employees of Open Source Development Network (OSDN), a corporation that sells advertising on Slashdot and customized implementations of the Slash platform. Stories are filtered out if they have technical formatting problems or, in principle, if they are poorly written or outdated. This segment of the service, then, seems mostly traditional—paid employees of the “publisher” decide which stories are, and which are not, interesting and of sufficient quality. The only peer production element here is the fact that the initial trolling of the web for interesting stories is itself performed in a distributed fashion. This characterization nonetheless must be tempered, because the filter is relatively coarse, as exemplified by the FAQ response to the question, “How do you verify the accuracy of Slashdot stories?” A Slashdot editor replied:

We don’t. You do. [] If something seems outrageous, we might look for some corroboration, but as a rule, we regard this as the responsibility of the submitter and the audience. This is why it’s important to read comments. You might find something that refutes, or supports, the story in the main.<sup>51</sup>

In other words, Slashdot is organized very self-consciously as a means of facilitating peer production of accreditation; it is at the comments stage that the story undergoes its most important form of accreditation—peer review ex post.

And things do get a lot more interesting as one looks at the comments. Slashdot allows the production of commentary on a peer-based model. Users submit comments that are displayed together with the initial submission of a story. Think of the “content” produced in these comments as a cross between academic peer review of journal submissions and a peer-produced substitute for television’s “talking heads.” It is in the means of accrediting and evaluating these comments that Slashdot’s system provides a comprehensive example of peer production of relevance and accreditation.

---

<sup>51</sup>. Open Source Dev. Network, Inc., *Slashdot FAQ: Editorial* (Oct. 28, 2000), at <http://slashdot.org/faq/editorial.shtml>.

2002]

## Coase's Penguin

Slashdot implements an automated system to select moderators from the pool of users.<sup>52</sup> Moderators are selected according to several criteria: They must be logged in (not anonymous), they must be regular users (average users, not one-time page loaders or compulsive users), they must have been using the site for a while (this defeats people who try to sign up just to moderate), they must be willing, and they must have positive “karma.” Karma is a number assigned to a user that primarily reflects whether the user has posted good or bad comments (according to ratings from other moderators). If a user meets these criteria, the program assigns the user moderator status and the user gets five “influence points” to review comments. The moderator rates a comment of his choice using a drop down list with words such as “flamebait” and “informative.” A positive word increases the rating of a comment one point and a negative word decreases the rating one point. Each time a moderator rates a comment, it costs the moderator one influence point, so the moderator can only rate five comments for each moderating period, which lasts for three days. If the user does not use the influence points within the period, they expire. The moderation setup is designed to give many users a small amount of power—thus decreasing the effect of rogue users or users with poor judgment. The site also implements some automated “troll filters,” which prevent users from sabotaging the system. The troll filters prevent users from posting more than once every sixty seconds, prevent identical posts, and will ban a user for twenty-four hours if the user has been moderated down several times within a short time frame.

Slashdot provides the users with a “threshold” filter that allows each user to block lower-quality comments. The scheme uses a numerical rating of the comment ranging from -1 to 5. Comments start out at 0 for anonymous posters, 1 for registered users, and 2 for registered users with good karma. As a result, if a user sets their filter at 1, the user will not see any comments from anonymous posters unless the comment’s ratings were increased by a moderator. A user can set their filter anywhere from -1 (viewing all of the comments) to 5 (viewing only the comments that have been upgraded by several moderators).

Users also receive accreditation through their karma. If their posts consistently receive positive ratings, their karma will increase. At a certain karma level, their comments will start off with a rating of 2, thereby giving them a louder voice, because other users with a threshold of 2 will now see their posts immediately, and fewer upward moderations are needed to push their comments even higher. Conversely, a user with bad karma from

---

52. The description in the following few paragraphs is mostly taken from the site’s frequently asked questions page or from observations. See Open Source Dev. Network, Inc., *Slashdot FAQ: Comments and Moderation*, at <http://slashdot.org/faq/com-mod.shtml> (last visited Sept. 1, 2002).

consistently poor ratings can lose accreditation by having their posts initially start off at 0 or -1. At the -1 level, the posts may not get moderated, effectively removing the opportunity for the “bad” poster to regain karma.

Relevance, as distinct from accreditation, is also tied into the Slashdot scheme because off-topic posts should receive an “off-topic” rating by the moderators and sink below the threshold level (assuming the user has the threshold set above the minimum). However, the moderation system is limited to choices that sometimes are not mutually exclusive. For instance, a moderator may have to choose between “funny” (+1) and “off topic” (-1) when a post is both funny and off topic. As a result, an irrelevant post can increase in ranking and rise above the threshold level because it is funny or informative. It is unclear whether this is a limitation on relevance or rather mimics our own normal behavior, say in reading a newspaper or browsing a library, where we might let our eyes linger longer on a funny or informative tidbit even after we’ve ascertained that it is not exactly relevant to what we were looking for.

In addition to mechanizing means of selecting moderators and minimizing their power to skew the aggregate judgment of the accreditation system, Slashdot implements a system of peer review accreditation for the moderators themselves. Slashdot implements this “meta-moderation” by making any user who has an account from the first ninety percent of accounts created on the system eligible to moderate the moderations. Each eligible user who opts to perform meta-moderation review is provided with ten random opportunities to rate moderators. The randomness helps to prevent biases and control by anyone who might use the assignment process to influence the selection of moderators. The user/meta-moderator may rate the moderator as either unfair, fair, or neither. The meta-moderation process affects the karma of the original moderator, which will remove the moderator from the moderation system, if lowered sufficiently.

Together, these mechanisms allow for the distributed production of both relevance and accreditation. Because there are many moderators who can moderate any given comment and mechanisms that explicitly limit the power of any one moderator to overinfluence the aggregate judgment, the system evens out differences in evaluation by aggregating judgments. The system then allows individual users to determine what level of accreditation fits their particular time and needs by setting their filter to be more or less inclusive. By introducing karma, the system also allows users to build reputation over time and to gain greater control over the accreditation of their own work relative to the power of the critics. Users, moderators, and meta-moderators are all volunteers. Slashdot demonstrates that the same dynamic that we observed for peer production of content can be implemented to produce relevance and accreditation. Rather than using the full-time effort of professional accreditation experts, the system is designed

2002]

Coase's Penguin

to permit the aggregation of many small judgments, each of which entails a trivial effort for the contributor. The software that mediates communication among the collaborating peers also contains a variety of mechanisms designed to defend the common effort from poor judgment or defection.

### C. *Value-Added Distribution*

After considering content production along with relevance and accreditation mechanisms, there remains the question of "distribution." To some extent this is a nonissue on the Internet. Distribution is cheap; all one needs is a server and large pipes connecting the server to the world, and anyone, anywhere, can get the information. I mention it here for two reasons. First, there are a variety of value-adding activities that need to be done at the distribution stage—like proofreading in print publication. Although the author who placed the content on the Web will likely, for the same motivations that caused him or her to put the materials together in the first place, seek to ensure these distribution values, we have very good examples of value-adding activities at the distribution stage being produced on a peer production model. Second, as the Internet is developing, the largest Internet Service Providers (ISPs) are trying to differentiate their services by providing certain distribution-related values. The most obvious examples are caching and mirroring—implementations by the ISP (caching) or a third party like Akamai (mirroring) that insert themselves into the distribution chain in order to make some material more easily accessible than other material.<sup>53</sup> The question is the extent to which peer distribution can provide similar or substitute values.

The most notorious example is Napster.<sup>54</sup> The collective availability of tens of millions of hard drives of individual users provided a substantially more efficient distribution system for a much wider variety of songs than

---

53. Part of the time lag involved in downloading materials is the time it takes for the materials to traverse the network from their point of origin to the user's computer. One approach to speeding up communications is to store copies of popular materials close to users. When Internet Service Providers do this, the function is called "caching," which relates to temporary storage of recently viewed files. See David D. Clark & Marjorie Blumenthal, *Rethinking the Design of the Internet: The End to End Arguments vs. the Brave New World*, Paper Presented at the Policy Implications of End-to-End Workshop, Stanford University, Dec. 1, 2000, at 15 (Aug. 10, 2000), *at* <http://lawschool.stanford.edu/e2e/papers/TPRC-Clark-Blumenthal.pdf>. Akamai is a business that provides similar functionality, allowing content providers to purchase the functionality independently of the decisions of an ISP. See AKAMAI, TURBO-CHARGING DYNAMIC WEBSITES WITH AKAMAI EDGESUITE (2001), *at* [http://www.akamai.com/en/resources/pdf/Turbocharging\\_WP.pdf](http://www.akamai.com/en/resources/pdf/Turbocharging_WP.pdf). So, for example, if CNN wants to be served quickly, but AT&T Worldnet is not caching CNN, CNN can use the services of Akamai to "mirror" its site in many important local markets so that whoever accesses the materials will receive more rapid service.

54. See generally Salon Media Group, *at* <http://dir.salon.com/topics/napster/> (last visited Sept. 1, 2002) (collecting a variety of stories and explanations of the rise and fall of said dearly departed).

the centralized (and hence easier to control) distribution systems preferred by the recording industry. The point here is not to sing the praises of the dearly departed (as of this writing) Napster. Setting aside the issue of content ownership, efficient distribution could be offered by individuals for individuals. Instead of any one corporation putting funds into building and maintaining a large server, end-users opened part of their hard drives to make content available to others. Although Napster required a central addressing system to connect these hard drives, Gnutella and other emerging peer-to-peer networks do not.<sup>55</sup> This is not the place to go into the debate over whether Gnutella has its own limitations, be they scalability or free-riding.<sup>56</sup> The point is that there are both volunteers and commercial software companies involved in developing software intended to allow users to set up a peer-based distribution system that will be independent of the more commercially controlled distribution systems, operating from the edges of the network to its edges,<sup>57</sup> rather than through a controlled middle.<sup>58</sup>

Perhaps the most interesting, discrete, and puzzling (for anyone who dislikes proofreading) instantiation of a peer-based distribution function is Project Gutenberg and the site set up to support it, Distributed Proofreaders. Project Gutenberg<sup>59</sup> is a collaboration of hundreds of volunteers who scan in and correct books so that they are freely available in digital form. Currently, Project Gutenberg has amassed around 6300 public domain e-texts through the efforts of volunteers and makes the collection available to everyone for free.<sup>60</sup> The e-texts are offered in ASCII format, which is the lowest common denominator, making it possible to reach the widest audience, but Project Gutenberg does not discourage volunteers from also offering the e-texts in markup languages. It contains a search engine that

---

55. See Andy Oram, *Gnutella and Freenet Represent True Technological Innovation*, O'REILLY NETWORK, May 12, 2000, at 3, at <http://www.oreillynet.com/pub/a/network/2000/05/12/magazine/gnutella.html>.

56. See Eytan Adar & Bernardo A. Huberman, *Free Riding on Gnutella*, 5 FIRST MONDAY 1 (Oct. 2, 2000), at [http://www.firstmonday.dk/issues/issue5\\_10/adar/](http://www.firstmonday.dk/issues/issue5_10/adar/). But see Clay Shirky, *In Praise of Free Loaders*, O'REILLY NETWORK, Dec. 1, 2000, at [http://www.oreillynet.com/pub/a/p2p/2000/12/01/shirky\\_freeloading.html](http://www.oreillynet.com/pub/a/p2p/2000/12/01/shirky_freeloading.html) (contesting Adar and Huberman's argument).

57. See Clay Shirky, *Communities, Audiences, and Scale*, at [http://www.shirky.com/writings/community\\_scale.html](http://www.shirky.com/writings/community_scale.html) (last visited Nov. 4, 2002).

58. Eben Moglen has argued that peer distribution is dramatically better than proprietary distribution, because social familiarity allows people to better guess their friends' and acquaintances' preferences than a centralized distributor. If individuals are provided with the freedom to give their friends music or any form of utterance that they believe they will like, the information will arrive in the hands of most everyone who would want it within a very small number of steps. Eben Moglen, *The dotCommunist Manifesto: How Culture Became Property and What We're Going To Do About It*, Lecture at the University of North Carolina, Chapel Hill (Nov. 8, 2001), at <http://www.ibiblio.org/moglen> (video stream).

59. Project Gutenberg, at <http://promo.net/pg/> (last visited Sept. 25, 2002).

60. See Ibiblio.org, at <http://www.ibiblio.org/gutenberg/GUTINDEX.ALL> (last visited Sept. 24, 2002) (including multiple versions of the same books as separate e-texts).

2002]

Coase's Penguin

allows a reader to search for typical fields, such as subject, author, and title. Distributed Proofreaders is a site that supports Project Gutenberg by allowing volunteers to proofread an e-text by comparing it to scanned images of the original book. The site is maintained and administered by one person.

Project Gutenberg volunteers can select any book in the public domain to transform into an e-text. The volunteer submits a copy of the title page of the book to Michael Hart—who founded the project—for copyright research. The volunteer is notified to proceed if the book passes the copyright clearance. The decision on which book to convert to e-text is thus left up to the volunteer, subject to copyright limitations. Typically, a volunteer converts a book to ASCII format using OCR (optical character recognition) and proofreads it one time in order to screen it for major errors. The volunteer then passes the ASCII file to a volunteer proofreader. This exchange is orchestrated with very little supervision. The volunteers use a listserv mailing list and a bulletin board to initiate and supervise the exchange. In addition, books are labeled with a version number indicating how many times they have been proofed. The site encourages volunteers to proof books that have low numbers. The Project Gutenberg proofing process is simple and involves looking at the text itself and examining it for errors. The proofreaders (aside from the first pass) are not expected to have access to the book or scanned images, but merely review the e-text for self-evident errors.

Distributed Proofreaders,<sup>61</sup> a site unaffiliated with Project Gutenberg, is devoted to proofing Project Gutenberg e-texts more efficiently by distributing the volunteer proofreading function in smaller and more information-rich modules. In the Distributed Proofreaders process, scanned pages are stored on the site and volunteers are shown a scanned page and a page of the e-text simultaneously so that the volunteer can compare the e-text to the original page. Because of the fine-grained modularity, proofreaders can proof one or a few pages and submit them. By contrast, the entire book is typically exchanged on the Project Gutenberg site, or at minimum a chapter. In this fashion, Distributed Proofreaders clears the proofing of thousands of pages every month.

Interestingly, these sites show that even the most painstaking, some might say mundane, jobs can be produced on a distributed model. Here the motivation problem may be particularly salient, but it appears that a combination of bibliophilia and community ties suffices (both sites are much smaller and more tightly knit than, for example, the Linux kernel development community). Individuals can self-identify as having a passion

---

61. Distributed Proofreaders, at <http://charlz.dns2go.com/gutenberg/> (last visited Aug. 31, 2002).

for a particular book or as having the time and inclination to proofread as part of a broader project they perceive to be in the public good. By connecting a very large number of people to these potential opportunities to produce, the e-text projects, just like Clickworkers, Slashdot, or Amazon, can capitalize on an enormous pool of underutilized intelligent human creativity and willingness to engage in intellectual effort.

#### D. *Summary*

What I hope these examples provide is a common set of mental pictures of peer production. In the remainder of the Article, I will abstract from these stories some general observations about peer production: what makes it work and what makes it better under certain circumstances than market- or hierarchy-based production. But at this point it is important that the stories have established the plausibility of, or piqued your interest in, the claim that peer production is an existing phenomenon of much wider application than free software. What remains is the interesting and difficult task of explaining the phenomenon so as to begin to think about the policy implications of the emergence of this strange breed in the middle of our information economy. I will by no stretch of the imagination claim to have completed this task in the following pages, but I hope to identify some basic regularities and organizing conceptions that will be useful to anyone interested in pursuing the answer. Even if you do not buy a single word of my initial efforts to theorize the phenomenon, however, seeing these disparate phenomena as instances of a general emerging phenomenon in the organization of information production should present a rich topic of study for organization theorists, anthropologists, institutional economists, and business people interested in understanding new production models in a ubiquitously networked environment.

## II. WHY WOULD PEER PRODUCTION EMERGE IN A NETWORKED ENVIRONMENT?

### A. *Locating the Theoretical Space for Peer Production*

There are many places to attempt to provide a theoretical explanation of peer production. One option would be to focus on the organization-of-production literature that would be most sympathetic to the sustainability and productivity of peer production. This might include the literature regarding trust-based modes of organizing production<sup>62</sup> or literature that

---

62. See, e.g., Paul S. Adler, *Market, Hierarchy, and Trust: The Knowledge Economy and the Future of Capitalism*, 12 ORG. SCI. 215 (2001).

2002]

## Coase's Penguin

focuses on internal motivation and its role in knowledge production.<sup>63</sup> Perhaps it makes sense to focus on cultural or sociological characteristics of peer communities as a central explanation of peer production, starting with mainstream sociological and anthropological literature of gift-giving and reciprocity.<sup>64</sup> There are applications that are rather close both online<sup>65</sup> and offline,<sup>66</sup> as well as in economic analysis of organization.<sup>67</sup> The advantage of doing so would be that these approaches have rich and detailed analytic tools with which to analyze the phenomenon of peer production. The disadvantage is that these approaches are outside the mainstream of economic theory, which in turn looms large in discussions of law and policy. In this early study of the phenomenon of peer production, it seems

---

63. See, e.g., Margit Osterloh & Bruno S. Frey, *Motivation, Knowledge Transfer, and Organizational Forms*, 11 ORG. SCI. 538 (2000).

64. A very early exploration of gift exchange is Franz Boas, *The Social Organization and the Secret Societies of the Kwakiutl Indians*, in ANNUAL REPORT OF THE U.S. NATIONAL MUSEUM FOR 1895, at 311 (1897), available at <http://www.canadiana.org/cgi-bin/ECO/mtq?doc=14300>. The locus classicus is MARCEL MAUSS, *THE GIFT: THE FORM AND REASON FOR EXCHANGE IN ARCHAIC SOCIETIES* (W.D. Halls trans., W.W. Norton 1990) (1925). Valuable insight is also provided by Marshall D. Sahlins, *On the Sociology of Primitive Exchange*, in *THE RELEVANCE OF MODELS FOR SOCIAL ANTHROPOLOGY* 139 (Michael Banton ed., 1965) and CLAUDE LÉVI-STRAUSS, *THE ELEMENTARY STRUCTURES OF KINSHIP* (James Harle Bell et al. trans., Beacon Press 1969) (1967).

65. Indeed, this is central to Raymond's discussion of open source development, see RAYMOND, *supra* note 2, though it is not entirely clear that his description in fact fits the gift literature, given how distant and potentially disconnected the act of giving in open source communities is from the act of receiving.

66. In the offline world, the academic community has been described as thriving on shared social commitments to the pursuit of truth, progress, and open collaboration. Science in particular has been the subject of sociological analysis of a productive enterprise. Classics are BERNARD BARBER, *SCIENCE AND THE SOCIAL ORDER* (1952), WARREN O. HAGSTROM, *THE SCIENTIFIC COMMUNITY* (1965), and ROBERT K. MERTON, *THE SOCIOLOGY OF SCIENCE* (1973). Studies of gift exchange flow from the 1925 fountainhead of Marcel Mauss. See MAUSS, *supra* note 64. Work in this vein has followed both in anthropology and sociology. For a review of this literature and its application to current debates over patenting basic research, see Arti Kaur Rai, *Regulating Scientific Research: Intellectual Property Rights and the Norms of Science*, 94 NW. U. L. REV. 77 (1999).

67. This is not to say that there is no literature within economics that attempts to use the gift-exchange literature to study economic phenomena. Examples are George A. Akerlof, *Labor Contracts as Partial Gift Exchange*, 97 Q.J. ECON. 543 (1982) and Rachel E. Kranton, *Reciprocal Exchange: A Self-Sustaining System*, 86 AM. ECON. REV. 830 (1996). See also JANET TAI LANDA, *The Enigma of the Kula Ring*, in *TRUST, ETHNICITY, AND IDENTITY: BEYOND THE NEW INSTITUTIONAL ECONOMICS OF ETHNIC TRADING NETWORKS, CONTRACT LAW, AND GIFT-EXCHANGE* 141 (1994); Ernst Fehr et al., *When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets*, 16 J. LAB. ECON. 324 (1998).

The past few years have seen a particular emphasis on studying reciprocity itself as an economic phenomenon. DAN M. KAHAN, *THE LOGIC OF RECIPROCITY: A THEORY OF COLLECTIVE ACTION AND LAW* (forthcoming 2002); see also ERNST FEHR & KLAUS M. SCHMIDT, *THEORIES OF FAIRNESS AND RECIPROCITY: EVIDENCE AND ECONOMIC APPLICATIONS* (Inst. for Empirical Research in Econ., Working Paper No. 75, 2001), at <http://www.unizh.ch/cgi-bin/iew/pubdb2>; BRUNO S. FREY & STEPHAN MEIER, *PRO-SOCIAL BEHAVIOR, RECIPROCITY, OR BOTH?* (Inst. for Empirical Research in Econ., Working Paper No. 107, 2002), at <http://www.unizh.ch/cgi-bin/iew/pubdb2>; Ernst Fehr & Armin Falk, *Psychological Foundations of Incentives*, Schumpeter Lecture at the Annual Conference of the European Economic Association (Nov. 2001), at <http://www.unizh.ch/cgi-bin/iew/pubdb2>.

more important to establish its baseline plausibility as a sustainable and valuable mode of production within the most widely used relevant analytic framework than to offer a detailed explanation of its workings. Doing so should provide wider recognition of the policy implications and create a space for more methodologically diverse inquiries.

At the most general level, we can begin by looking at Ronald Coase's explanation of the firm in *The Nature of the Firm* and Harold Demsetz's explanation of property rights in *Toward a Theory of Property Rights*.<sup>68</sup> Coase's basic explanation of the emergence and relative prevalence of firms focuses on the comparative costs of institutional alternatives. In other words, Coase asked why clusters of individuals operate under the direction of an entrepreneur, a giver of commands, rather than interacting purely under the guidance of prices, and answered that using the price system is costly. Where the cost of achieving a given outcome in the world through the price system will be higher than the cost of using a firm to achieve the same result, firms will emerge. Any given firm will cease to grow when the increased complexity of its organization makes its internal decision costs higher than the costs that a smaller firm would incur to achieve the same marginal result. Firms will not, however, conduct activities if the cost of organizing these activities within a firm exceeds the cost of achieving that result through the market. Assuming that the cost of organization increases with size, Coase posited that we have a "natural"—i.e., internal to the theory—limit on the size and number of organizations.

Demsetz's basic explanation of why property emerges with regard to resources that previously were managed without property rights—as commons—is based on a very similar rationale. As long as the cost of implementing and enforcing property rights in a given resource is higher than the value of the increase in the efficiency of utilization of the resource gained by the introduction of a property regime, the resource will operate without property rights. An increase in the value of the resource due to an exogenous circumstance, such as a technological development or an encounter with another civilization, may create a sufficient incentive for property rights to emerge. More generally, property in a resource emerges if the social cost of having no property in that resource exceeds the social cost of implementing a property system in it. This restatement can include common property regimes, managed commons, and other nonproperty approaches to managing sustainable commons.<sup>69</sup>

Table 1 describes the interaction between Coase's theory of the firm and Demsetz's theory of property.

---

68. Harold Demsetz, *Toward a Theory of Property Rights*, 57 AM. ECON. REV. 347 (1967).

69. For discussions of commons, see OSTRÖM, *supra* note 17; and Rose, *supra* note 17. A brief discussion of these concepts as applied to peer production follows below. See *infra* notes 112-120 and accompanying text.

2002]

Coase's Penguin

TABLE 1. IDEAL ORGANIZATIONAL FORMS AS A FUNCTION OF  
RELATIVE SOCIAL COST

	Property system more valuable than implementation costs <sup>1</sup>	Implementation costs of property system higher than opportunity cost <sup>2</sup>
Market exchange of $x$ more efficient than organizing $x$	Markets	Commons
Organizing $x$ more efficient than market exchange of $x$	Firms	Common property regimes

<sup>1</sup> A property system is “valuable” as compared to the option, and opportunity costs, of not having property rights in place. The concept of the “value” of a property system in the first column is equivalent to the “opportunity cost” of a property system in the second column.

<sup>2</sup> Both markets and firms generally rely on property rights. Therefore, the institutions described in this column reflect functional equivalents for decentralized exchange—markets—and coordinating organization—firms—in the absence of property.

Before going into why peer production may be less costly than property/market-based production or organizational production, it is important to recognize that if we posit the existence of such a third option it is relatively easy to adapt the transaction-costs theory of the firm and the comparative institutional cost theory of property to include it. We would say that when the cost of organizing an activity on a peered basis is lower than the cost of using the market or hierarchical organization, then peer production will emerge.<sup>70</sup>

We could tabulate as follows:

TABLE 2. IDEAL ORGANIZATIONAL FORMS AS A FUNCTION OF  
RELATIVE SOCIAL COST INCLUDING PEER PRODUCTION

	Property system more valuable than implementation costs	Implementation costs of property system higher than opportunity cost
Market exchange of $x$	Markets	Commons

---

70. In the context of land, Ellickson extends Demsetz’s analysis in precisely this fashion, suggesting that there may be a variety of reasons supporting group ownership of larger tracts, including the definition of efficient boundaries, coping with significant shocks to the resource pool, risk spreading, and the “advent of inexpensive video cameras or other technologies for monitoring behavior within a group setting.” Robert C. Ellickson, *Property in Land*, 102 YALE L.J. 1315, 1330, 1332-44 (1993).

more efficient than organizing or peering of $x$	(Farmers markets)	(Ideas & facts; roads)
Organizing $x$ more efficient than market exchange or peering of $x$	Firms (Automobiles; shoes)	Common property regimes (Swiss pastures)
Peering of $x$ more efficient than organizing or market exchange of $x$	Proprietary “open source” efforts (Xerox’s Eureka)	Peer production processes <sup>1</sup> (NASA Clickworkers)

<sup>1</sup> “Cost” here would include the negative effects of intellectual property on dissemination and downstream productive use.

Understanding that the same framework that explains the emergence of property and firms could, in principle, also explain the emergence of peer production focuses our effort on trying to understand why it is that peering could, under certain circumstances, be a more cost-effective institutional form than either markets or hierarchical organizations. Because the emergence of peer production seems to be tied to the emergence of a pervasively networked information economy, my explanation seeks to be in some sense sensitive to (1) changes in the special characteristics of the human and material resources used in information production relative to other productive enterprises and (2) the cost and efficiency of communication among human participants in the productive enterprise.

#### B. *Peer Production of Information in a Pervasively Networked Environment*

Peer production is emerging as an important mode of information production because of four attributes of the pervasively networked information economy. First, the object of production—information—is quirky as an object of economic analysis, in that it is purely nonrival,<sup>71</sup> and its primary nonhuman input is the same public good as its output—information.<sup>72</sup> Second, the physical capital costs of information production have declined dramatically with the introduction of cheap processor-based computer networks. Third, the primary human input—creative talent—is

---

71. A good is nonrival to the extent that its consumption by one person does not diminish its availability for use by any other person. See, e.g., Paul M. Romer, *Endogenous Technological Change*, 98 J. POL. ECON. S71, S73-S74 (1990). It has been commonplace for a long time to treat information as a perfectly nonrival good. See *id.*; see also Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in NAT’L BUREAU OF ECON. RESEARCH, THE RATE AND DIRECTION OF INVENTIVE ACTIVITY: ECONOMIC AND SOCIAL FACTORS 609, 616-17 (1962).

72. While the input characteristic of information has been appreciated at least since 1962, Arrow, *supra* note 71, the extensive exploration of the implications of this characteristic largely begins with Suzanne Scotchmer, *Standing on the Shoulders of Giants: Cumulative Research and the Patent Law*, 5 J. ECON. PERSP. 29 (1991).

2002]

## Coase's Penguin

highly variable, more than traditional labor and certainly more than many material resources usually central to production. Moreover, the individuals who are the “input” possess better information than anyone else about the suitability of their talents and their level of motivation and focus at a given moment to given production tasks. Fourth, and finally, communication and information exchange across space and time are much cheaper and more efficient than ever before, which permits the coordination of widely distributed potential sources of creative effort and the aggregation of actual distributed effort into usable end products.

The first attribute—the public goods nature of information—affects the cost of one major input into production—existing information. It means that the social cost of using existing information as an input into new information production is zero.<sup>73</sup> This has two effects on the relative cost of peer production of information.<sup>74</sup> First, it lowers the expected private cost of peer production of *information*, as compared to normal economic goods, because in principle it means that a central input—preexisting information—could be available to human productive agents without limit, if the provisioning problem can be solved without introducing appropriation. Second, it underlies a pervasive social cost of market and hierarchy in this field of production because of the losses in both static and dynamic efficiency entailed by the implementation of property rights in a nonrival public good usually thought necessary to sustain market and hierarchy-based production of information.<sup>75</sup>

The second attribute—the decline in physical capital cost—similarly lowers the cost of another major cost of information production. The age of mechanical reproduction that enabled fixation and distribution of information and culture as goods was defined by the high cost of physical capital. Large circulation automated printing presses, vinyl record and later CD manufacturing facilities, and movie studios and their celluloid-based systems formed the basis for the industrial model typical of information and cultural production in the twentieth century. The declining cost of computer processors coupled with the digitization of all forms of information and

---

73. Saying that the social cost of its use by one person is zero is simply another way of saying that the good is nonrival—that its use by one person does not prevent its use by any other person.

74. The public goods aspect of information production is usually described as being comprised of two distinct characteristics, its nonrivalry and its nonexcludability. See Romer, *supra* note 71, at S73-S74. A good is excludable to the extent that its producer can exclude others from its use unless they pay. If a good is not excludable, it too presents a problem for market provisioning, not because it is inefficient to price it above zero, but because it is difficult to do so, and hence firms will provide too little of it. Nonexcludability of information is less important to the analysis here, because it does not relate to the characteristics of information that are important to making peer production both feasible and efficient—that its most efficient price is zero and that it can be used by any number of people without diminishing its availability for others.

75. See Arrow, *supra* note 71, at 617 (“[P]recisely to the extent that [property rights in information] are successful, there is an underutilization of the information.”).

culture has made the necessary physical capital cheaper by orders of magnitude than in the past.

Together, these first two attributes make information production a potentially sustainable low-cost, low-returns endeavor for many individuals relying on indirect appropriation.<sup>76</sup> It is important to note, however, that the public goods attribute limits the applicability of my observations about peer production, so that I make no claim about the applicability of these observations to traditional economic goods.

The third characteristic—the centrality of human capital to information production and its variability—is the primary source of efficiency gains from moving from markets or hierarchical organization to peering. Commons-based peer production creates better information about available human capital and can better allocate creative effort to resources and projects.

The fourth attribute—the dramatic decline in communications costs—radically reduces the cost of peering relative to its cost in the material world. This allows substantially cheaper movement of information inputs to human beings, human talent to resources, and modular contributions to projects, so that widely dispersed contributions can be integrated into finished information goods. It also allows communication among participants in peer production enterprises about who is doing what and what needs to be done.

### C. Markets, Hierarchies, and Peer Production as Information-Processing Systems

Usually, the question of why anyone would contribute to a peer production enterprise without directly appropriating the benefits is foremost in people's minds when I describe the phenomenon. For the sake of completeness of the organization-theory argument, however, suspend disbelief for one more section (or if you cannot do so, read Part III first, and then come back here). Assume for the next ten pages that I have come up with a reasonably plausible story as to why people participate, which will

---

76. "Indirect appropriation" is appropriation of the value of one's effort by means other than reliance on the excludability of the product of the effort. For example, someone who is paid as a teacher but gets the position in reliance on his scholarship is indirectly appropriating the benefit of his scholarship. An IBM engineer who gains human capital by working on GNU/Linux from home in the evening is indirectly appropriating the benefits of her efforts in participating in the production of GNU/Linux. The term is intended to separate out appropriation that is sensitive to excludability of information—direct appropriation through intellectual property—and appropriation that is independent of exclusion from the information—indirect appropriation without intellectual property. See Yochai Benkler, *Intellectual Property and the Organization of Information Production*, 22 INT'L REV. L. & ECON. 81, 87-88 (2002). As a general matter, the more a sector of information production can be sustained through indirect appropriation, the less it needs intellectual property.

2002]

## Coase's Penguin

allow us to consider the claim that if they did, their efforts would be more productive than if they organized in a firm or interacted purely through a price system.

Peer production has a relative advantage over firm- or market-based production along two dimensions, both a function of the highly variable nature of human capital. The first emerges when one treats all approaches to organizing production as mechanisms by which individual agents reduce uncertainty as to the likely value of various courses of productive action.<sup>77</sup> Differences among these modes in terms of their information-processing characteristics could then account for differences in their relative value as mechanisms for organizing production. The second dimension consists of allocation efficiencies gained from the absence of property. A particular strategy that firms, and to a lesser extent markets, use to reduce uncertainty is securing access to limited sets of agents and resources through contract and property. This strategy entails a systematic loss of allocation efficiency relative to peer production, because there are increasing returns to scale for the size of the sets of agents and resources available to be applied to projects and peer production relies on unbounded access of agents to resources and projects.

### 1. *Information Gains*

We could reduce the decisions that must be made by productive human beings as follows. Imagine a human agent (*A*), who is deciding whether and how to act. An action (*a*) is a combination of two elements: the effort (*e*) to be exercised and the resources (*r*) as to which the effort is exercised. An agent sees a range of actions open to him that is a function of the range of resources and effort levels that he can utilize. As information-collection costs decline, agents see more of the universe of opportunities for action that are available to them. Imagine that *A* is a rational actor,<sup>78</sup> where the

---

77. Individuals who are presented with alternatives from which to choose, such as standing in a particular spot and turning a lever all day or writing an economic analysis of friendship, do not always know which of these courses of action is more valuable or which would allow them to put dinner on the table. One can think of markets and firms as means by which individuals solve this lack of knowledge, because they use signals about which action will better fulfill their purposes—be they glory or subsistence. What follows, then, is in some measure a sketchy application of Herbert Simon's statement: "It is only because individual human beings are limited in knowledge, foresight, skill, and time that organizations are useful instruments for the achievement of human purpose . . ." HERBERT A. SIMON, MODELS OF MAN 199 (1957). This individual-centric view of organization treats the firm solely in terms related to the question of why agents use this form of organization to order their individual productive behavior. I do not differentiate between entrepreneurs, managers, and employees, but rather treat all of them as agents who have a set of possible open courses of action.

78. A rational actor in the most traditional sense can be formalized within a framework that strictly orders the value of outcomes of actions performed by agents who can calculate the values of their preferences, outcomes, and the probabilities of outcomes vis-à-vis actions. For the limited

private value ( $V_A$ ) to  $A$  of doing  $a$  is the expected value of the outcome ( $O_A$ ), which is the value to  $A$  of  $O_A$  obtaining, discounted by the probability that  $O_A$  will obtain if  $A$  does  $a$ . This means that the value to  $A$  of doing  $a$  increases as the probability that doing  $a$  will result in  $O_A$  obtaining increases. That probability depends on the effort  $A$  will exert; the resources available to  $A$ ;  $A$ 's talent ( $t$ ), where talent describes relative capabilities, associations, and idiosyncratic insights and educational mixes of an individual that make that person more or less productive with a given set of resources for a given project; the presence of complementary actions by other agents; and the absence of undermining actions by other agents.

$A$  will do  $a_n$  using  $(e_n, r_n)$  if  $A$  believes the value  $V_{An}$  to be higher than either inaction or an alternative action. This requires that  $V_{An}$  be positive relative to the value of inaction and higher than the value  $V_{Am}$ —the value of any other  $O_{Am}$  similarly discounted by the probability that any other  $a_m$ , combining any  $(e_m, r_m)$ , will lead to  $O_{Am}$  obtaining. It is important to underscore that the probability of  $O_A$  obtaining is in some measure dependent on the actions, both complementary and undermining, of other agents. Assuming that the agent knows his or her own valuation of the outcome, has some experience-based evaluation of  $t$ , and controls his or her own effort,  $e$ , uncertainty resides primarily with regard to the divergence of the private valuation of the outcome from its valuation by others, the availability of  $r$ , and the interdependence of the agent's action on the action or inaction of others outside  $A$ 's control. Reducing these uncertainties is a central function of markets and firms. Reducing the latter two in particular is a central function of property and contract, which can secure complementary material and human resources to increase the probability that complementary actions will be taken and decrease the probability that undermining actions will occur.

Markets and firm-based hierarchies are information processes in the sense that they are means of reducing the uncertainty that agents face in evaluating different courses of action to a level acceptable to the agent as warranting action. Markets price different levels of effort and resources to signal the relative values of actions in a way that allows individuals to compare actions and calculate the likely actions of other individuals faced with similar pricing of alternative courses of action. Firms reduce uncertainty by specifying to some individuals what actions to take, thereby reducing the uncertainty of interdependent action to a level acceptable to

---

purposes of comparing the information-processing characteristics of firms, markets, and peer production, these rather strong characteristics are not necessary. It is enough to have individuals that, in Simon's terms, are satisficers. *See id.* at 204-05. The uncertainty as to the relative value of a given action must be reduced to a level that satisfies the actor's requirements to justify action, without needing to calculate fully the various outcomes.

2002]

## Coase's Penguin

the agents by delegating to the managers control over enough resources and people by contract and property.

To compare modes of organizing production as information-processing systems one might use the term *information opportunity cost*. I use the term “information” here in the technical sense of a reduction in uncertainty, where “perfect information” is the condition where uncertainty regarding an action could not be further reduced in principle. Perfect information is impossible to acquire, and different organizational modes have different strategies for overcoming this persistent uncertainty. These strategies differ from each other in the amount and kind of information they lose in the process of resolving the uncertainty that rational agents face. The divergence of each mode from the hypothetical condition of perfect information—its lossiness—is that mode’s information opportunity cost.

Markets reduce uncertainty regarding allocation decisions by producing a signal that is clear and comparable across different uses as to which use of the relevant factors would be most efficient. To do so, markets require a codification of the attributes of different levels of effort, different kinds of resources, and different attributes of outcomes, so that these can all be specified as contract terms to which a price is affixed. An example of this was the introduction of codified standards for commodities as an indispensable element of the emergence of commodities markets.<sup>79</sup>

Since we are concerned with individual agents’ decisions, and levels and focuses of effort are a major component of individual action, it is intuitive that specification and pricing of all aspects of individual effort as they change in small increments over the span of an individual’s full day, let alone a month, is impossible.<sup>80</sup> What we get instead is codification of effort types—a garbage collector, a law professor—that are priced more or less finely. But one need only look at the relative homogeneity of law firm starting salaries as compared to the high variability of individual ability and motivation levels of graduating law students to realize that pricing of individual effort can be quite crude. These attributes are also difficult to monitor and verify over time, though perhaps not quite as difficult as predicting them *ex ante*, so that pricing continues to be a function of relatively crude information about the actual variability among people.

---

79. See ALFRED D. CHANDLER, JR., *THE VISIBLE HAND: THE MANAGERIAL REVOLUTION IN AMERICAN BUSINESS* 211 (1977) (describing how commodity attributes became codified, and local variability squelched, as a part of the transition to commodity markets).

80. In the context of the market for labor, this has sometimes been called the multitask problem—the inability to specify completely by contract all the tasks required and all the attributes of an employee who will likely need to perform multiple tasks. Bengt Holmstrom & Paul Milgrom, *Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design*, 7 J.L. ECON. & ORG. 24 (1991) (Special Issue); see also Bengt Holmstrom, *The Firm as a Subeconomy*, 15 J.L. ECON. & ORG. 74, 89 (1999) (“The fact that [job] contingencies are hard to specify *ex ante* makes the firm a potentially important operator of an internal human capital market.”).

More importantly, as aspects of performance that are harder fully to specify in advance or monitor—like creativity over time given the occurrence of new opportunities to be creative—become more important, market mechanisms become more lossy.

Markets are particularly good at resolving the uncertainties with regard to the difference in valuation of the outcome among different agents. Therefore, an agent acting on a market price will have a relatively certain evaluation of the external valuation of the outcome. Of course, this valuation may be flawed because of externalities not reflected in the price, but (for better or worse, depending on the magnitude and shape of externalities) markets plainly reduce uncertainty about the value of an action as perceived by others. Markets also reduce the uncertainty about the availability of resources, by allowing an agent to compare the value of an outcome to the price of necessary resources. Finally, markets reduce uncertainties with regard to the actions of other agents in two ways. First, agents can evaluate the risk that others will act in a way that is detrimental, or fail to act in a way that is complementary, to the agent's action, given the relative pricing of the courses of complementary or detrimental action. This risk assessment can then be built into the perceived value of a possible action. Second, agents can maintain property rights in resources and projects to prevent negatively correlated actions<sup>81</sup> and contract with other individuals to provide relatively secure access to resources for complementary action.

Firms or hierarchical organizations resolve uncertainty by instituting an algorithm or process by which information about which actions to follow is ordered, so that some pieces of information count as a sufficient reduction in uncertainty about the correct course of action to lead agents who receive them to act. The mythical entrepreneur (or the historical manager) becomes the sole source of information that is relevant to reducing the uncertainty of the workers in a purely managed firm. In the ideal-type firm, uncertainty as to which of a set of actions will increase an agent's welfare is reduced by fixing a salary for following a manager's orders and shifting some of the risk of that course of action from employees to employers. Production processes—if I stand here and twist this lever all day, cars will emerge from the other side and I will get a paycheck—are codified as instruction sets. Agents reduce their uncertainty about why to act and what to do by reducing the universe of information they deem relevant to their decision.

---

81. Maintaining rights in what I call “projects” is, on Kitch’s now-classic reading, the primary function of the patent system. See Edmund Kitch, *The Nature and Function of the Patent System*, 20 J.L. & ECON. 265 (1977). Even if one is critical of Kitch’s almost exclusive focus on this characteristic as the reason for the patent system, recognizing that patents provide some measure of control over projects is all that is necessary here. The derivative use right in copyright plays a similar function to a more limited extent.

2002]

## Coase's Penguin

Information that arrives through a particular channel with a particular level of authorization counts as a signal, and all the rest counts as noise. It remains to the entrepreneur (in the pure model of the firm)<sup>82</sup> to be the interface between the firm and the market and to translate one set of uncertainty-reducing signals—prices—to another set of signals with similar effect—organizational commands.

By controlling a set of resources through property and commanding a set of agents through the employment relationship, the firm reduces the elements of uncertainty related to the interdependence of the actions of agents. But by doing so it creates a boundary around the set of available agents and the set of available resources that limits the information available about what other agents could have done with these same resources or what else these agents could have done with other resources. This boundary therefore limits the efficacy of information-collection mechanisms—like incentive-based contracts—that firms use to overcome the difficulty of collecting information to which their employees have special access.<sup>83</sup>

The point to see is that like the price system, hierarchical organization is a lossy medium. All the information that could have been relevant to the decision regarding each factor of production but that was not introduced in a form or at a location that entitled it to “count” toward an agent’s decision is lost. Much of the knowledge-management movement in business schools and punditry since the mid-1990s has been concerned with mitigating the lossiness of managerial hierarchy as an information-processing mechanism.<sup>84</sup> Mitigating this lossiness is the primary job of CIOs.<sup>85</sup>

---

82. A number of readers have complained that this picture of the firm is too thin to be a realistic and complete description. Firms use all sorts of market-based mechanisms like incentive compensation and internal arm’s length bargaining among units and mix market-based and hierarchical control mechanisms to organize production. The point of my description, however, is not to present a true sociological description of production in a firm. Coase’s transaction costs theory of the firm identifies two dimensions to the process of allocating resources—pricing and managerial commands—making it possible to map different organizations according to whether and how they mix these ideal-type modes of coordinating the use of resources in production. I present these two ideal types in their ideal form here, so as to clarify what is different about peer production. Within this theory, peer production emerges as a third ideal type, to create a three-dimensional space within which an organizational model can be described. In this model, employee-of-the-month programs and employee-feedback sessions become simple precursors to hybrids between firms and peer production processes, the most obvious example of which is presented by Xerox’s Eureka.

83. Another problem with incentive-based contracts is that they may undermine voluntary cooperation, a phenomenon related to the relationship between the presence of money and social-psychological rewards discussed *infra* notes 101-105. See ERNST FEHR & SIMON GÄCHTER, DO INCENTIVE CONTRACTS UNDERMINE VOLUNTARY COOPERATION? (Inst. for Empirical Research in Econ., Working Paper No. 34, 2002), at <http://www.unizh.ch/cgi-bin/iew/pubdb2>.

84. For a range of definitions of knowledge management and a taste of the analytic styles, see Brint.com, What Is Knowledge Management?, at <http://www.brint.com/km/kmdefs.htm> (last visited Aug. 31, 2002).

An example where peer production—proprietary, not commons-based—was used precisely to solve the lossiness of hierarchical organization is Xerox's Eureka system for organizing the flow of questions from, and answers to, field technicians about failures of photocopiers.<sup>86</sup> The firm created a decentralized communications system for technicians to post questions, a peer review system for other technicians to answer these questions, and a database library of past questions and answers available to technicians who confront new problems. The original approach toward technical failures of machines was to use manuals that came with the machines, because the machine was conceptualized as being completely engineered by the engineers, with all the possible failures specified in the manual. Technicians were thus conceived of as instruction followers, who came to machines that were broken, diagnosed the problem by locating it in the manual, and then solved it by executing a series of corrective steps prescribed by the manual.

The Eureka project changed the conception of the knowledge content of the machine and the organizational role of technicians from instruction followers to knowledge producers. In this system, a technician who came across a problem not clearly resolved in the manual posted a question electronically to a proprietary communications system accessible by all technicians. Any other technician in the system who had come across a similar problem could post a fix, which would be reviewed by experienced technicians, who would opine on its advisability. The technician who used the fix could then report on whether it worked. The whole transaction was stored in a database of solutions. The technicians were not compensated for answering queries but instituted instead a system of authorship and honor-based payoffs. Eureka also flipped the traditional hierarchical conception of knowledge in a machine as codified by engineers and implemented by instruction-following technicians. The knowledge content of the machine was now understood to be something that is incomplete when it leaves the design board and is completed over the life of the machines by technicians who share questions and solutions on a peer-review, volunteer model.

The Eureka project suggests one additional interpretation of peer production in relation to markets and hierarchies. We generally understand the existence of markets and hierarchies as two ideal models of organization and observe various mixes of the two types in actual organizational practices. Eureka suggests that peer production can be a third ideal-type organizational model, which can be combined in various

---

85. CIOs are “Chief Information Officers,” a position created to reduce information loss within organizations.

86. See Daniel G. Bobrow et al., *Community Knowledge Sharing in Practice*, at <http://jonescenter.wharton.upenn.edu/VirtualCommunities/whalen.pdf> (last visited Aug. 31, 2002).

2002]

## Coase's Penguin

measures with the other two, forming a three-dimensional map of organizational strategies rather than the two-dimensional map recognized traditionally.

Recognizing the lossiness of markets and managerial hierarchies suggests the first portion of a working hypothesis about why peer production has succeeded in gaining ground, namely the possibility that peer production may have lower information opportunity costs than markets or hierarchies. In particular, I suggest that the primary source of gains—which may be called *information gains*—that peer production offers is its capacity to collect and process information about human capital. The hypothesis is that rich information exchange among large sets of agents free to communicate and use existing information resources cheaply will create sufficiently substantial information gains that, together with the allocation gains that I will discuss in the following Section, overcome the information-exchange costs due to the absence of pricing and managerial direction and the added coordination costs created by the lack of property and contract.

Where the physical capital costs of information production are low and where existing information resources are freely or cheaply available,<sup>87</sup> the low cost of communication among very large sets of agents allows agents to collect information through extensive communication and feedback instead of using information-compression mechanisms like prices or managerial instructions. If communications include a sufficiently large number of agents operating in the same resource and opportunity sets, this mode of communication can provide to each agent rich information about what needs to or can be done, who is doing what, and how other people value any given outcome. One sees this phenomenon in the centrality of effective communications platforms to the design of peer production processes—be they simple lists that lie at the heart of every free software development project<sup>88</sup> or the sophisticated collaboration platforms that underlie projects like Slashdot or Kuro5hin. The value of these systems is precisely in enabling agents to use extensive information exchange and feedback to provide the same desiderata that prices and managerial commands provide in their respective models.

---

87. By limiting the hypothesis to information production under conditions of cheap and widely available physical capital (computers) and relatively free availability of information inputs, we can largely ignore uncertainty as to the availability of material resources, because the domain of application of the hypothesis relates to conditions where resources other than human creativity are not scarce, so that uncertainty as to their availability is minimal.

88. At the heart of the distributed production system that is typified by open source software development is the notion of making the program available in a publicly accessible space for people to comment on and upgrade. See RAYMOND, *supra* note 9, at 26-28. These communication lists have also offered a valuable location for observers of the phenomenon to gain insight. See, e.g., Karim R. Lakhani & Eric von Hippel, *How Open Source Software Works: "Free" User-to-User Assistance*, 32 RES. POL'Y (forthcoming 2003) (describing support lists for Apache), available at <http://opensource.mit.edu/papers/lakhanivonhippelusersupport.pdf>.

Platform design and maintenance, and more importantly the human attention required to take in and use this information, are the equivalent for peer production of organization and decision costs in firms and of transaction costs in markets. The magnitude of these costs will partly be a function of the quality of the design of the collaboration platform in terms of efficiency of communication and information-processing utilities. This rich information exchange may or may not be efficient, depending on the magnitude of the cost and the relative information gains generated by the richer information available to agents through this system.

Reducing uncertainty about the availability of opportunities for action by any given agent and about complementary actions by other agents becomes the salient potential source of information gain for peer production projects, while the capacity of a project to reduce the likely prevalence or efficacy of undermining actions becomes a major limiting factor. This latter effect, most obviously typified by the information-rich process of peer review, will occupy a substantial portion of Part III, where I will discuss in some detail the threats to effective peer production and the mechanisms available to this mode of production to defend itself from incompetence and defection. Here I will focus on the information gains generated by peer production in terms of opportunities both for creative and for novel utilization of existing resources<sup>89</sup> and opportunities agents have to use their own talents, availability, focus, and motivation to perform a productive act.

Central to my hypothesis about the information gains of peer production is the claim that human intellectual effort is highly variable and individuated. People have different innate capabilities, personal, social, and educational histories, emotional frameworks, and ongoing lived experiences. These characteristics make for immensely diverse associations with, idiosyncratic insights into, and divergent utilization of, existing information and cultural inputs at different times and in different contexts. Human creativity is therefore very difficult to standardize and specify in the contracts necessary for either market-cleared or hierarchically organized production. As human intellectual effort increases in importance as an input into a given production process, an organization model that does not require contractual specification of effort but allows individuals to self-identify for tasks will be better at gathering and utilizing information about who should be doing what than a system that does require such specification. Intra-firm hybrids, like incentive compensation, may be able to improve on firm-only or market-only approaches, but it is unclear how well they can overcome

---

89. This is a point Bessen makes about complex software, *see* Bessen, *supra* note 12, as well as a characteristic of the motivation Raymond describes as having an itch to scratch, RAYMOND, *supra* note 9, at 23 (“Every good work of software starts by scratching a developer’s personal itch.”).

2002]

Coase's Penguin

the core difficulty of requiring significant specification of the object of organization and pricing—in this case, human intellectual input.

The point here is qualitative. It is not only, or even primarily, that more people can participate in production. *The widely distributed model of information production will better identify who is the best person to produce a specific component of a project, all abilities and availability to work on the specific module within a specific time frame considered.* With enough uncertainty as to the value of various productive activities and enough variability in the quality of information inputs and human creative talent vis-à-vis any set of production opportunities, coordination and continuous communications among the pool of potential producers and consumers can generate better information about the most valuable productive actions and the best human agents available at a given time. Although markets and firm incentive schemes are aimed at producing precisely this form of self-identification, the rigidities associated with collecting and comprehending bids from individuals through these systems (i.e., transaction costs) limit the efficacy of self-identification, relative to peer production.

Now, self-identification is not always perfect, and some mechanisms used by firms and markets to codify effort levels and abilities—like formal credentials—are the result of experience with substantial errors or misstatements by individuals about their capacities. To succeed, therefore, peer production systems must also incorporate mechanisms for smoothing out incorrect self-assessments, like peer review in traditional academic research or the major sites like Slashdot or Kuro5hin, or like redundancy and statistical averaging in the case of NASA Clickworkers. In information terms, these mechanisms reduce the uncertainty associated with the likely presence of undermining actions by other agents. The prevalence of misperceptions that agents have about their own abilities and the cost of eliminating such errors will be part of the transaction costs associated with this form of organization that are parallel to quality control problems faced by firms and markets. This problem is less important where the advantage of peer production is in acquiring fine-grained information about motivation and availability of individuals who have otherwise widely available capabilities—like the ability to evaluate the quality of someone else's comment on Slashdot. It is likely more important where a particular skill set is necessary that may not be widespread—like the programming skills necessary to fix a bug in a program.

## 2. Allocation Gains

In addition to its potential information gains, peer production has potential allocation gains enabled by the large sets of resources, agents, and

projects available to peer production. This gain is cumulative to the general information-processing characteristics of peer production and is based on the high variability of human capital, which suggests that there are increasing returns to the scale of the pool of individuals, resources, and projects to which they can be applied.

As illustrated in Figure 1(a), market- and firm-based production processes rely on property and contract to secure access to bounded sets of agents and resources in the pursuit of specified projects. The permeability of the boundaries of these sets is limited by the costs of making decisions in a firm about adding or subtracting a marginal resource, agent, or product and the transaction costs of doing any of these things through the market. Peer production relies on making an unbounded set of resources available to an unbounded set of agents, who can apply themselves toward an unbounded set of projects. The variability in talent and other idiosyncratic characteristics of individuals suggests that any given resource will be more or less productively used by any given individual and that the overall productivity of a set of agents and a set of resources will increase when the size of the sets increases toward completely unbounded availability of all agents to all resources for all projects. Even if in principle the decisionmaker has information as to who is the best person for a job given any particular set of resources and projects (in other words, if the information gains are assumed away), the transaction or organizational costs involved in bringing that agent to bear on the project may be too great relative to the efficiency gain over use of the resource by the next-best available agent who is within the boundary.

Assume that the productivity ( $P_A$ ) of a set of agents/resources is a function of the agents ( $A$ ) available to invest effort ( $e$ ) on resources ( $r$ ). The productivity of Agent  $A_1$  ( $P_{A1}$ )<sup>90</sup> is a function of the set of resources on which  $A_1$  can work ( $r_1$ ), the level of effort  $A_1$  will invest ( $e_1$ ), and  $A_1$ 's talent ( $t_1$ ).  $P_A$  increases as a function of  $e$ ,  $r$ , and the actions of other complementary agents and decreases by undermining actions of agents at a magnitude that is a function of  $t$ . Note that  $t$  is a personal characteristic of individuals that is independent of the set of resources open for  $A$  to work on, but will make a particular  $A$  more or less likely to be productive with a given set of resources in collaboration with other agents for the achievement of a set of outcomes.  $A$ 's access to any given set of resources and potential collaborators therefore represents a probability, which is a

---

90. While  $V_A$  discussed in the previous Section related to the private value of an action to an agent,  $P_A$  is intended to represent the potential social value of the efforts of any one or more agents  $A$  as part of a potential collaborative effort.

2002]

## Coase's Penguin

function of  $t$ , that  $A$  will be productive with that resource and those collaborators for a given project.<sup>91</sup>

The existence of  $t$  generates increasing returns to scale of the set of resources and to the set of agents to which it is available, because the larger the number of agents with access to a larger number of resources, the higher the probability that the agents will include someone with a particularly high  $t$  for productive use of a given  $r_n$  at a given  $e_n$  as compared to other agents. Imagine a scenario where  $A_1$  works for Firm  $F_1$  and has a higher  $t$  value as regards using  $r_2$  than  $A_2$  who works for Firm  $F_2$ . If  $r_2$  is owned by  $F_2$ ,  $r_2$  will be used by the less efficient  $A_2$ , so long as the value of  $A_2$  working on  $r_2$  is no less than the value of  $A_1$  working on  $r_2$  minus the transaction costs involved in identifying the relative advantage of  $A_1$  and assigning  $A_1$  to work on  $r_2$ . This potential efficiency loss would be eliminated if  $A_1$  were in the set of agents who had transaction-cost-free access to work on the resource set that includes  $r_2$ .

If Firms  $F_1$  and  $F_2$  each have a set of agents and resources,  $\{A_{F1}, r_{F1}\}$  and  $\{A_{F2}, r_{F2}\}$ , then  $P_{F1} + P_{F2} < P_{F1+F2}$ .

Figures 1(a) and 1(b) illustrate the point with another example. Figure 1(a) assumes that there are two firms, each having contracts with a set of agents and property in a set of resources. Assume that as among  $\{A_1 \dots A_5\}$  in Firm  $F_1$  the best agent for using the combination  $(r_1, r_4)$  is  $A_2$ . Assume also that as among the agents  $\{A_1 \dots A_9\}$ ,  $A_8$  is the best, in the sense that if  $A_8$  were to use these resources, the social value of the product would be greater by some measure ( $m$ ) than when  $A_2$ , the best agent within Firm  $F_1$ , uses them.

It is unlikely that the two firms will have the information that  $A_8$  is best for the job, as I suggested in the discussion of information gains. Even if they do know, creativity will be misapplied as long as the transaction costs associated with transferring the creativity of  $A_8$  to Firm  $F_1$  or the property in  $r_1$  and  $r_4$  to Firm  $F_2$  are greater than  $m$ . When the firms merge, or when the agents and resources are in a commons-based peer production enterprise space, the best person can self-identify to use the resources, as in Figure 1(b). Think of this as someone musing about fairy tales and coming up with a biting satire, which she is then capable of implementing, whereas the employee of the initial owner of the rights to the fairy tale might only produce a depressingly earnest new version.

---

91. In seeking to identify the private value of an outcome to an agent above, I described the successful completion of a project as an outcome  $O_A$ , discounted by the probability  $q$  of  $O_A$  obtaining should  $A$  do  $a$ .  $P_A$  is the social equivalent of  $qO_A$  to the individual, representing a judgment of whether an individual will be productive.


This initial statement is a simplification and understatement of the potential value of the function by which the sizes of the sets of agents and resources increase productivity. There are two additional components: the range of projects that might be pursued with different talent applied to a given set of resources and the potential for valuable collaboration. First, a more diverse set of talents looking at a set of resources may reveal available projects that would not be apparent when one only considers the set of resources as usable by a bounded set of agents. In other words, one of the advantages may be not the ability of  $A_1$  to pursue a given project with  $r_2$  better than  $A_2$  could have but the ability to see that a more valuable project is possible. Second, the initial statement does not take into consideration the possible ways in which cooperating individuals can make each other creative in different ways than they otherwise would have been. Once one takes into consideration these diverse effects on the increased possibilities for relationships among individuals and between individuals and resources, it becomes even more likely that there are increasing returns to scale to increases in the number of agents and resources involved in a production process.

2002]


Coase's Penguin

FIGURE 1. APPLYING AGENTS TO RESOURCES

(a) Separated in Different Firms



(b) In a Common Enterprise Space



Assume, for example, that every agent, given a  $t$  value, has some potential ability to use every resource, which could be measured as an option for that agent on that resource. (In other words, its value is derived from the value of the agent using the resource well, discounted by the probability that the agent will be good at using the resource.) Assume also that every agent has some potential to add value in collaboration with any other agent, and that every resource could have some potential value in combination with any other resource. If we have one agent,  $A_1$ , and one resource,  $r_1$ , we only get the value of the option of  $A_1$  using  $r_1$ . If we add one more resource, we get the value of  $A_1$  using  $r_1$ ,  $A_1$  using  $r_2$ , and  $A_1$  using  $r_1$  in combination with  $r_2$ . Symmetrically, if we keep the resource set fixed at one resource but add an agent, because the resources are nonrival, we would see the value of two agents and one resource as the sum of the values of  $A_1$  using  $r_1$ ,  $A_2$  using  $r_1$ , and a collaboration between  $A_1$  and  $A_2$  to use  $r_1$ . If we combine adding one agent and one resource, we see the following. The value of  $\{A_1, r_1\} + \{A_2, r_2\}$ , if the two sets are strictly separated, is the value of  $A_1$  using  $r_1$  and  $A_2$  using  $r_2$ . The value of  $\{A_1, A_2, r_1, r_2\}$  in a single agent/resource space is the combined value of  $A_1$  using  $r_1$  and  $A_2$  using  $r_2$ ,  $A_1$  using  $r_2$  and  $A_2$  using  $r_1$ ,  $A_1$  using  $r_1$  and  $r_2$  and  $A_2$  using  $r_1$  and  $r_2$ ,  $A_1$  and  $A_2$  collaborating to use  $r_1$ ,  $A_1$  and  $A_2$  collaborating to use  $r_2$ , and  $A_1$  and  $A_2$  collaborating to use  $r_1$  and  $r_2$ .<sup>92</sup>

Figure 2 illustrates this point. Each arrow identifies one potential option for a valuable combination of agents and resources. In Figure 2(a), we see that separating the two agents and resources results in a combined value of only two options. Figure 2(b) shows the three combinations of a single agent with two resources. Likewise, Figure 2(c) represents the three options associated with two agents and one resource. In Figure 2(d), we see that, once both agents and resources are placed in the same opportunity set, the number of options for use and collaboration increases dramatically, with each arrow representing one of the nine potentially valuable combinations of agents and resources that the single agent-resource space makes possible.

---


92. I am not sure there is room to formalize the precise relationship here on the style of Metcalfe's Law or Reed's Law. See David P. Reed, That Sneaky Exponential—Beyond Metcalfe's Law to the Power of Community Building, at <http://www.reed.com/Papers/GFN/reedslaw.html> (last visited Aug. 31, 2002). From a policy perspective, there is no need to do so at this early stage of studying the phenomenon. It is sufficient for our purposes here to see that the collaboration effects and insights due to exposure to additional resources mean that the returns to scale are, as with other networks, more than proportional.

2002]


Coase's Penguin

FIGURE 2. AGENT AND RESOURCE COMBINATIONS


(a) Bounded Sets of One Agent and One Resource




(b) One Agent Combined with Two Resources



(c) Two Agents Combined with One Resource



(d) Two Agents Combined with Two Resources



2002]

Coase's Penguin

If this is true, then in principle a world in which all agents can act effectively on all resources will be substantially more productive in creating information goods than a world in which firms divide the universe of agents and resources into bounded sets. As peer production relies on opening up access to resources for a relatively unbounded set of agents, freeing them to define and pursue an unbounded set of projects that are the best outcome of combining a particular individual or set of individuals with a particular set of resources, this open set of agents is likely to be more productive than the same set could have been if divided into bounded sets in firms. Note that the effect changes dramatically when the resources are rival, because the value of any agent or combination of agents working on the resource is not additive to the value of any other agent or combination. In other words, the use of a rival resource excludes the use by others in a way that is not true for a purely nonrival good like information. The allocation gain is attained in allocating the scarce resource—human attention, talent, and effort—given the presence of nonrival resources to which the scarce resource is applied with only a probability of productivity.

This is not to say that peer production will always necessarily be more productive or that it will always improve with size. First, adding agents may increase coordination and communication costs and heighten the probability that the set of agents will include individuals whose actions, through incompetence or malice, will undermine the productivity of the set.

Second, in situations where focused but relatively standardized effort is more important than variability of talent, well-understood incentive systems based on monetary rewards could outweigh this effect, and markets and firms are much better understood mechanisms to generate the incentives for such an application of effort than is peer production. Even in these situations, however, monetary incentives will not necessarily be more efficient, even if better understood. If a project can be structured to resolve the effort/incentives problem without appropriation of the output, as I describe in Part III, the substantial increases in productivity resulting from the availability of a larger set of resources to a larger set of agents with widely variable talent endowments could be enough to make even an imperfectly motivated peer production process more productive than firms that more directly motivate effort but segment agents and resources into smaller bounded sets. Moreover, as Section III.A explains, a peer production project could increase, rather than decrease, motivation by eliciting contributions motivated by nonmonetary rewards when monetary rewards would have been either ineffective or inefficient.

And third, as unbounded sets of agents utilize unbounded sets of resources for unbounded sets of projects, there is likely to be substantial duplication of effort. This duplication is wasteful if one considers actual likely patterns of peer production as compared to an idealized peer

production system where everyone self-identifies perfectly for the one contribution that they are best suited to produce. The question then becomes one of comparative efficiency: How much of a drag is duplication on the claimed increased efficiency provided by peer production enterprises? The answer has two primary components. First, as Part III will elaborate, peer production draws effort that in many cases would otherwise have been directed toward purely nonproductive consumption—say, watching television instead of marking craters on Mars, ranking websites for the Open Directory Project, or authoring entries for Wikipedia. On a macrolevel of social productivity, then, it enables an economic system to harness activities from which agents gain welfare to peer production efforts that also generate innovation and welfare for others. While duplication may limit the total value of this newly tapped source of productivity, it is less important to the extent that the duplication occurs among efforts that would, in the absence of a peer production system, have gone unused in the production system. Second, and probably more important, redundancy provides important values in terms of the robustness and innovativeness of a system. Having different people produce the same component makes the production system more resistant to occasional failures. Moreover, having different people with different experience and creative approaches attack the same problem will likely lead to an evolutionary model of innovation where alternative solutions present themselves, thus giving the peer production process the ability to select among a variety of actual solutions rather than precommitting to a single solution.

### III. OF MOTIVATION AND ORGANIZATION: THE COMMONS PROBLEM

#### A. *The “Incentives” Problem: Of Diverse Motivations and Small Contributions*

What makes contributors to peer production enterprises tick? Why do they contribute? There are two versions of this question. The first is the question of the economic skeptic. It questions the long-term sustainability of this phenomenon, given that people will not, after the novelty wears off, continue to work on projects in which they can claim no proprietary rights.<sup>93</sup> It is to this question that my discussion here responds, in an effort to show that the network as a whole can be a sustainable system for the production of information and culture. There is a second, narrower version

---

93. This skepticism is more often encountered in questions in conferences and presentations than in formal papers. A well-articulated written example of a skeptic’s view, however, is Glass, *supra* note 11, comparing recruiting operating system developers to Tom Sawyer’s whitewashing the fence trick and arguing that eventually operating system efforts will die because too many important programming tasks are not fun/sexy enough.

2002]

## Coase's Penguin

of the question, which arises once one overcomes the skepticism and begins to consider how peer production can be steered or predicted. It would seek to understand the motivations and patterns of clustering around projects in the absence of property rights and contracts and the emergence of the effective networks of peers necessary to make a *particular* project succeed. These are questions that present rich grounds both for theoretical and empirical study. My hunch is that these would best be done informally in the domains of social psychology and anthropology, or, if done formally, through artificial life-type modeling. They are, in any event, beyond the scope of this initial study, which is intended solely to define the phenomenon and assess its sustainability and welfare effects in general terms.

As a practical matter, the incentive problem as an objection to the general sustainability of peer production is in large part resolved by the existence of a series of mechanisms for indirect appropriation of the benefits of participation catalogued quite comprehensively by Lerner and Tirole.<sup>94</sup> At the broadest level, there is the pleasure of creation. Whether you refer to this pleasure dispassionately as “hedonic gain” or romantically as “an urge to create,” the mechanism is simple. People are creative beings. They will play at creation if given an opportunity, and the network and free access to information resources provide this opportunity.<sup>95</sup> More closely related to the project of keeping body and soul together, there is a variety of indirect appropriation mechanisms for those who engage in free software development. These range from the amorphous category of reputation

---

94. Lerner & Tirole, *supra* note 10.

95. Moglen makes this central to his explanation in *Anarchism Triumphant* as follows:

It's an emergent property of connected human minds that they create things for one another's pleasure and to conquer their uneasy sense of being too alone. The only question to ask is, what's the resistance of the network? Moglen's Metaphorical Corollary to Ohm's Law states that the resistance of the network is directly proportional to the field strength of the “intellectual property” system.

....  
So, in the end, my dwarfish friends, it's just a human thing. Rather like why Figaro sings, why Mozart wrote the music for him to sing to, and why we all make up new words: Because we can. Homo ludens, meet Homo faber.

Moglen, *supra* note 21. Raymond, as well as Lerner and Tirole, also offer hedonic gains as one component of their respective explanations. RAYMOND, *supra* note 2, at 13-14; Lerner & Tirole, *supra* note 10, at 213. There is, of course, something counterintuitive about calling hedonic pleasure “indirect” appropriation. I use the terms “direct” and “indirect” to distinguish between appropriation that relies directly on the economic exclusion made possible by intellectual property law and all other forms of appropriation. Direct appropriation supports the claim that intellectual property rights can increase information production, while indirect appropriation undermines the claim. On why it is that these two types of appropriation have this oppositional relationship as insights into the utility of the intellectual property rights, see Benkler, *supra* note 76 (explaining that intellectual property has positive effects on information-production strategies that are based on direct appropriation and negative effects on information-production strategies that rely on indirect appropriation, and that these effects will structure the organization of information production and its efficiency even when they have no effect on aggregate productivity in information production).

gains<sup>96</sup> to much more mundane benefits, such as consulting contracts, customization services, and increases in human capital that are paid for by employers who can use the skills gained from participation in free software development in proprietary projects. In this regard, it is important to note that about two-thirds of the revenues in the software industry are not tied to software publishing, but to service-type relationships.<sup>97</sup> Given that two-thirds of the revenues of the software industry are service-based and that the total revenues of the software industry are three times the size of the movie, video, and sound-recording industries combined,<sup>98</sup> indirect appropriation offers a rich field of enterprise for participants in free software development.

The reality of phenomena like academic research, free software, the World Wide Web, NASA Clickworkers, and Slashdot supports these explanations with robust, if not quantified here, empirical grounding. All one need do is look at the Red Hat founders (no longer billionaires, but not quite on the bread line either)<sup>99</sup> and IBM's billion-dollar commitment to

---

96. But, as mentioned above, this commonly cited motivation has not been reconciled with the contrary practices of two of the most successful free software projects, Apache and the Free Software Foundation, neither of which provides personal attribution to code they bless. *See supra* note 7.

97. The Economic Census of 1997 breaks up software into several categories, ranging from publishing to different types of services and education. Nonetheless, it is possible to collect the information about the industry as a whole using primarily the following three categories: software publishing (NAICS 5112), computer systems design and related services (NAICS 5415), and computer training (NAICS 61142). *See U.S. CENSUS BUREAU, ECONOMIC CENSUS OF 1997, at* <http://www.census.gov/epcd/ec97/us/US000.HTM> [hereinafter ECONOMIC CENSUS OF 1997]. Software publishing had receipts of more than \$61 billion (35%), computer systems design and related services roughly \$109 billion (63%), and computer training roughly \$2.5 billion (1.4%). *Id.*

The Economic Census of 1997 divides computer systems design and related services into a number of subcategories: custom computer programming services (NAICS 541511, \$38 billion); computer systems design services (NAICS 541512, \$51 billion); computer facilities management services (NAICS 541513, \$15 billion); and other computer related services (NAICS 541519, \$4.3 billion). *Id.* Finally, computer systems design services is further subdivided into computer systems integrators (NAICS 5415121, \$35 billion) and computer systems consultants (NAICS 5415122, \$16 billion). *Id.* The Economic Census defined computer systems design and related services as follows:

This industry comprises establishments primarily engaged in providing expertise in the field of information technologies through one or more of the following activities: (1) writing, modifying, testing, and supporting software to meet the needs of a particular customer; (2) planning and designing computer systems that integrate computer hardware, software, and communication technologies; (3) on-site management and operation of clients' computer systems and/or data processing facilities; and (4) other professional and technical computer-related advice and services.

U.S. Census Bureau, NAICS 54151: Computer Systems Design and Related Services, *at* <http://www.census.gov/epcd/ec97/def/54151.HTM> (last visited Oct. 27, 2002).

98. The movie, video, and recording industries (NAICS 512) had total receipts of roughly \$56 billion, as compared to roughly \$173 billion for the software industry. ECONOMIC CENSUS OF 1997, *supra* note 97.

99. Red Hat is a company that specializes in packaging and servicing GNU/Linux operating systems. In 1999, it had an immensely successful IPO that made its founders billionaires, for a

2002]

Coase's Penguin

supporting Linux and Apache, on the one hand, and the tens of thousands of volunteer clickworkers, thousands of Linux developers, and hundreds of distributed proofreaders, on the other hand, to accept intuitively that some combination of hedonic gain and indirect appropriation can resolve the incentives problem. In this Part, I abstract from this intuitive observation to offer an answer that is more analytically tractable and usable to understand the microanalytic questions of peer production and the potential range in which peer production will be more productive than firms or markets.

### 1. *Abstracting the Effect of Diverse Motivations*

Saying that people participate for all sorts of reasons is obviously true at an intuitive level. It does not, however, go very far toward providing a basis for understanding why some projects draw many people, while others fail, or how the presence or absence of money affects the dynamic. What I will try to do in this Subsection is propose a framework to generalize the conditions under which peer production processes will better motivate human effort than market-based enterprises. Given the discussion of the information and allocation gains offered by peer production, this Subsection outlines a range in which peer production should be more productive than market-based or firm-based production. At the broadest level, wherever peer production can motivate behavior *better* than markets or firms, then certainly it will be superior. It will also be potentially better over a range where it may motivate behavior less effectively than markets or firms, but the contribution of the lower overall effort level will be less than the contribution of the added value in terms of information about, and allocation of, human creativity.

Let any agent have a set of preferences for rewards of three types:

$M \rightarrow$  *Monetary rewards*, which decrease in value because of the decreasing marginal utility of money. Call the rate at which  $M$  decreases  $s$  (satiation).

$H \rightarrow$  *Intrinsic hedonic rewards* experienced from taking the actions.

$SP \rightarrow$  *Social-psychological rewards*, which are a function of the cultural meaning associated with the act and may take the form of actual effect on social associations and status perception by others or on internal satisfaction from one's

---

while. The company, however, survived the bursting of the dot-com bubble and continues to lead the field of Linux distributions.

social relations or the culturally determined meaning of one's action.

At an intuitive level, three common examples help to clarify this diversity of motivation. Simplest to see is how these motivations play out with regard to sex: the prostitution fee (*M*), the orgasm (*H*), and love (*SP*). One can also make and serve dinner to others for any combination of a fee (*M*), the pleasure of cooking (*H*), and companionship (*SP*). The combination of these interacting motivations shapes our understanding of whether we are observing a short-order cook, a restaurant chef, or a dinner party host. Similarly, one can write about law for a legal fee (*M*), the pleasure of creating a well-crafted argument (*H*), or the respect of the legal community or one's colleagues (*SP*). To some extent, all three exist for anyone writing, but in different measures partly depending on taste and partly depending on social role, such as whether the author is a practitioner, a judge, or an academic, as well as on other factors, such as external time constraints.

The value of the three types of rewards for any given action might be independent of the value of the others, or it might not.<sup>100</sup> For purposes of this analysis, I will assume that *H* is a personal preference that is independent of the other two<sup>101</sup> but that *M* and *SP* can be positively or negatively correlated depending on the social construction of having money associated with the activity. I will call this factor *p*, which can be negative (as in prostitution) or positive (as in professional sports).<sup>102</sup> The *p* factor is most interesting when it is negative and is intended to allow for the possibility of a "crowding-out" phenomenon,<sup>103</sup> which has mostly been

---

100. Needless to say, the independent value of each may be positive or negative. One might be willing to pay money to engage in hedonically pleasing or social-psychologically satisfying activities, as people do all the time for hobbies, and people often take hedonically unpleasant, socially awkward, or even demeaning jobs in order to get the positive monetary rewards.

101. Separating out purely physical pleasure or pain from the social-psychological meaning of the cause of the pleasure or pain is artificial in the extreme. In principle, hedonic gain can be treated as part of *SP*, and indeed I ignore it as an independent factor in the analysis. I include it in the general statement largely to separate out the social-psychological aspect, which, unlike hedonic gains, is usually downplayed in economics.

102. Again, the culturally contingent nature of the relationship should be obvious. When the Olympics were renewed in the modern era, they were limited to "amateurs," because professional sports were a form of entertainment, giving their paid performers no more respect than paid performers were given more generally. As with all performers, this changed with the status inversion that was part of the twentieth-century celebrity culture generated to focus mass demand on mass-produced entertainment as opposed to the relationship/presence-based entertainment of the past.

103. See Bruno S. Frey & Reto Jegen, *Motivation Crowding Theory: A Survey of Empirical Evidence*, 15 J. ECON. SURV. 589 (2001); Bruno S. Frey & Felix Oberholzer-Gee, *The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out*, 87 AM. ECON. REV. 746 (1997); see also FEHR & GÄCHTER, *supra* note 83 (analyzing the ability of incentive contracts to undermine voluntary cooperation). For a broader moral claim about this tradeoff, see MARGARET

2002]

## Coase's Penguin

studied in the context of the relatively rare instances where altruistic provisioning has been the major, if not exclusive, mode of provisioning of socially important material goods among strangers, such as blood<sup>104</sup> or gametes.<sup>105</sup> While analysis leaves serious questions as to whether altruistic provisioning of these types of goods is indeed superior to market-based provisioning as a general social policy,<sup>106</sup> the primary disagreement concerns which mode is more efficient in the aggregate, not whether market provisioning displaces altruistic provisioning or whether each mode draws different contributors.<sup>107</sup> Using our three intuitive examples, an act of love drastically changes meaning when one person offers the other money at its end, and a dinner party guest who takes out a checkbook at the end of dinner instead of bringing flowers or a bottle of wine at the beginning likely never will be invited again. The question of money in legal writing will depend on the social construction of the role of the author. For a practicing advocate, *p* usually is positive, and higher monetary rewards represent the respect the author receives for her craft. For a judge, *p* with regard to payment for any particular piece of writing is strongly negative, representing the prohibition on bribes. For academics, *p* for a particular piece of writing may be positive or negative, depending on whether its

---

JANE RADIN, CONTESTED COMMODITIES (1996), and for a critique, see Kenneth J. Arrow, *Invaluable Goods*, 35 J. ECON. LITERATURE 757 (1997).

104. The quintessential source of the claim that altruism is superior to markets in providing blood is RICHARD M. TITMUSS, THE GIFT RELATIONSHIP: FROM HUMAN BLOOD TO SOCIAL POLICY (1971).

105. See Ken R. Daniels, *Semen Donors: Their Motivations and Attitudes to Their Offspring*, 7 J. REPROD. INFANT PSYCHOL. 121 (1989) (finding that a majority of donors surveyed in Australia and New Zealand gave altruistic reasons as the major motivation, with payment rated as rather unimportant); see also Simone B. Novaes, *Giving, Receiving, Repaying: Gamete Donors and Donor Policies in Reproductive Medicine*, 5 INT'L J. TECH. ASSESSMENT HEALTH CARE 639 (1989) (reviewing motivation for, and social issues of, sperm and egg donation and surrogacy); Simone B. Novaes, *Semen Banking and Artificial Insemination by Donor in France: Social and Medical Discourse*, 2 INT'L J. TECH. ASSESSMENT HEALTH CARE 219 (1986) (providing an account of several sperm banks in France and an analysis of their varying policies on donor compensation). But see Linda S. Fidell et al., *Paternity by Proxy: Artificial Insemination with Donor Sperm*, in GENDER IN TRANSITION: A NEW FRONTIER 93, 100 (Joan Offerman-Zuckerberg ed., 1989) (reporting that three-quarters of sperm donors in the United States were primarily motivated by financial gain).

106. Titmuss's thesis was challenged in a series of papers in the 1970s, see, e.g., Kenneth J. Arrow, *Gifts and Exchanges*, 1 PHIL. & PUB. AFF. 343 (1972); Robert M. Solow, *Blood and Thunder*, 80 YALE L.J. 1696 (1971), and more recently has been subject to refinement with the experience of the AIDS epidemic, see Kieran Healy, *The Emergence of HIV in the U.S. Blood Supply: Organizations, Obligations, and the Management of Uncertainty*, 28 THEORY & SOC'Y 529 (1999).

107. Specifically for an evaluation of Titmuss's argument in light of the HIV crisis, see Kieran Healy, *Embedded Altruism: Blood Collection Regimes and the European Union's Donor Population*, 105 AM. J. SOC. 1633, 1637-54 (2000) (reporting on an international comparison and concluding that "the opportunity to sell plasma does reduce one's likelihood of giving blood"). More generally, for a description of empirical surveys in a number of areas, see Frey & Jegen, *supra* note 103 (describing empirical research in multiple disciplines supporting the displacement effect money has on voluntaristic motivations).

source is considered to be an interested party paying for something that is more akin to a brief than to an academic analysis, or, for example, a foundation or a peer-reviewed grant, in which case “winning” the support is considered as adding prestige.

A distinct motivational effect arises when  $SP$  is associated with participation in collective action and concerns the presence or absence of rewards to the other participants and the pattern of the reward function—that is, whether some people get paid and others do not or whether people get paid differentially for participating. This relationship could be positive where altruism or a robust theory of desert culturally structures the social-psychological component of the reward to support monetary appropriation by others<sup>108</sup> or, more commonly perhaps, negative where one agent is jealous of the rewards of another. I will denote this factor  $jalt$  (jealousy/altruism).

Agents will then face different courses of action that they will perceive as having different expected rewards  $R$ :

$$R = M_s + H + SP_{p,jalt}$$

At any given time, an agent will face a set of possible courses of action and will have a set of beliefs about the rewards for each course of action, each with this form. A rational agent will choose based on the value of  $R$ , not of  $M$ . Irrespective of one’s view of whether the agent is a maximizer or a satisficer, the agent will have some total valuation of the rewards for adopting differing courses of action and hence of the opportunity cost of following courses of action that exclude other courses of action.

It is quite intuitive to see then that there will be some courses of action whose reward will be heavily based on hedonic or social-psychological parameters, on primarily monetary rewards, or on a combination of all three factors. At the broadest level one can simply say that agents will take actions that have a positive value and low opportunity cost because they do not displace more rewarding activities. Similarly, where opportunities for action do compete with each other, an agent will pursue an activity that has low, zero, or even negative monetary rewards when the total reward, given the hedonic and social-psychological rewards, is higher than alternative courses of action that do have positive monetary rewards attached to them. Hence the phenomena of starving artists who believe they are remaining true to their art rather than commercializing or of law professors who

---

108. A religiously motivated agent, for example, might consider the acquisition of monetary returns by other agents a positive sign of success, because the appropriators are seen as deserving in whatever theory of desert is prescribed by the religion, such as neediness or having been chosen in some sense.

2002]

Coase's Penguin

forego large law-firm partner draws when they choose teaching and writing over the practice of law.

What more can we say about the likely actions of agents whose preferences for rewards take the form I describe? First, there is a category of courses of action that will only be followed, if at all, by people who seek social-psychological and hedonic rewards. Assume that there are transaction costs for defining and making  $M$  and  $SP$  available to the agent,  $C_m$  and  $C_{sp}$ , respectively. I assume that these costs are different, because the former require the definition and enforcement of property rights, contracts, and pricing mechanisms, while the latter require social mechanisms for the association of social-psychological meaning with the act generally and with the individual agent's act in particular.

There is potentially a category of cases where the marginal value ( $V$ ) of an agent's action will be less than the transaction costs of providing monetary rewards for it, in which case the expected monetary reward will be zero. If the social value of the contribution is greater than zero, however, and if the hedonic and social-psychological rewards are greater than zero and greater than the cost of making the social-psychological rewards available, then it will be socially efficient for agents to act in this way when opportunities to act arise. Agents will in fact do so if someone has incurred the costs of providing the opportunities for action and the social-psychological or hedonic rewards.

Behaviors in the following range will therefore occur only if they can be organized in a form that does not require monetary incentives and captures behaviors motivated by social-psychological and hedonic rewards:

$$C_m > V > C_{sp} \text{ and } H + SP - C_{sp} > 0$$

Whether this range of activities is important depends first on the granularity of useful actions. The more fine-grained the actions and the more of these small-scale actions that need to be combined into a usable product, the higher the transaction costs of monetizing them relative to the marginal contribution of each action.<sup>109</sup>

Second, approaches that rely on social-psychological rewards will be particularly valuable to motivate actions that are systematically undervalued in the market, because they generate high positive externalities. A fairly intuitive example is basic science, which is particularly ill-suited for proprietary information production because of its high positive externalities,<sup>110</sup> and in which our social-cultural framework has developed

---

109. Technology that lowers the transaction costs could counteract this effect, decreasing the size of the group of cases that fall into this category.

110. The public goods problem of information production limits the efficacy of proprietary provisioning under any circumstances. The fact that basic science has many and varied uses as a

an elaborate honor-based rewards system rather than one focused on monetary rewards. We see similar social-psychological reward structures to reward and motivate participation in other practices that produce high positive externalities that would be difficult fully to compensate in monetary terms, like teaching; military service; or uncorrupt political, cultural, or spiritual leadership. Similarly, to the extent that peer production can harness motivations that do not require monetization of the contribution, the information produced using this model can be released freely, avoiding the inefficiencies associated with the public goods problem of information.

It is important to recognize that actions involved in creating the opportunities for others to act are themselves acts with analogous reward structures. The scientists who created the Mars Clickworkers project operated on one set of monetary and social-psychological returns, while the clickworkers themselves operated in response to a different set of hedonic and social-psychological returns. The Open Source Development Network funds the Slashdot platform based on one set of rewards, including an expected monetary return, but its action generates opportunities for others to act purely on *SP*- and *H*-type rewards. The crucial point is that the presence of *M*-type rewards for the agent generating the opportunities does not necessarily negatively affect the social-psychological returns to agents who act on these opportunities. In other words, there is some reason why the *jalt* factor for the contributors is not a strong negative value even though monetary factors are captured by the person providing the opportunity for collaboration.

We need, then, to state the relationship between the presence of *M*-type rewards for an action and the *SP*-type rewards associated with it. For simplicity, I will treat the total effect of both modifiers of *SP* as *p*, and will separate out *jalt* only where there is a reason to differentiate between the effect of monetary returns to the agent and the effects of differential reward functions for different agents in a collaborative group—as in the case where the person offering the opportunity to collaborate has different rewards from the participants in the collaboration.

---

fundamental input into new innovation and learning creates particularly large positive externalities and makes the public goods problem particularly salient in that context. The point was initially made in Nelson, *supra* note 26, at 306 (“The problem of getting enough resources to flow into basic research is basically the classical external-economy problem.”). The point was reiterated in Arrow, *supra* note 71, at 623-25. A particularly helpful detailed discussion is Rebecca S. Eisenberg, *Intellectual Property at the Public-Private Divide: The Case of Large-Scale cDNA Sequencing*, 3 U. CHI. L. SCH. ROUNDTABLE 557 (1996). For more general statements of the relationship between academic and commercial work, see Richard R. Nelson, *What Is “Commercial” and What Is “Public” About Technology, and What Should Be?*, in *TECHNOLOGY AND THE WEALTH OF NATIONS* 57, 65-70 (Nathan Rosenberg et al. eds., 1992). See also Ralph Gomory, *The Technology-Product Relationship: Early and Late Stages*, in *TECHNOLOGY AND THE WEALTH OF NATIONS*, *supra*, at 383, 388.

2002]

Coase's Penguin

Keeping hedonic gains to one side, the reward function can be represented as:

$$R = M_s + SP_p$$

This function suggests two key implications. First, we can confidently say that whenever  $M$  and  $SP$  are independent of each other or are positively correlated (that is, when  $p \geq 0$ ), approaches that provide monetary rewards for an activity will be preferred to nonmonetizable approaches toward the exact same activity. A rational agent will prefer a project that provides both social-psychological and monetary rewards over one that offers only one of these rewards. Someone who loves to play basketball will, all other things being equal, prefer to be paid for playing at Madison Square Garden over playing at West Third and Sixth Avenue without being paid.

Second, we can say that when  $M$  and  $SP$  are negatively correlated ( $p < 0$ ), an activity will be more or less attractive to agents depending on the values of  $s$  and  $p$ —that is, on the rate at which the value of marginal monetary rewards for a new action is discounted by the agent and the rate at which the presence of money in the transaction devalues the social-psychological reward for that action. Table 3 maps the likely effects of monetary rewards on the value of  $R$  as a function of the values of  $s$  and  $p$ . We can say generally that individuals with a high discount rate on money (high  $s$ ) will be likely to pursue activities with a high absolute value negative  $p$  rate only if these are organized in a nonproprietary model, because the value of  $M_s$  for them is low, and the presence of any  $M$ -type reward substantially lowers the value of  $SP_p$ . At the simplest level, this could describe relatively wealthy people—for example, a wealthy person is unlikely to take a paying job serving lunch at a soup kitchen, but may volunteer for the same job. More generally, most people who have finished their day job and are in a part of their day that they have chosen to treat as leisure, even though a second job is available, can be treated as having a higher  $s$  value for that part of their day. During this portion of the day, it will likely be easier to attract people to a project with social-psychological benefits, and if  $p$  is large and negative, adding monetary rewards will lower, rather than increase, participation. As we move toward a situation where the value of  $s$  for an individual is low, and the  $p$  rate, though negative, is low, we will tend to see a preference for combining  $M$  and  $SP$ , as one would where  $p$  is neutral or positive.

TABLE 3. EFFECTS OF INCREASED MONETARY REWARDS ON  $R$ 

		<i>S</i>	
		High	Low
$ p $ , where $p < 0$	High	Substantial decrease of $R$	Substantial or insignificant increase or decrease of $R^1$
	Low	Insignificant increase or decrease of $R$	Substantial increase of $R$

<sup>1</sup> An increase in monetary rewards in this quadrant may be substantial or insignificant in either increasing or decreasing  $R$ , because an increase in monetary rewards will substantially decrease  $SP$  but will substantially increase  $M$ . If the substantial decrease in  $SP$  and increase in  $M$  are roughly equivalent, then the overall result will be an insignificant change of  $R$ . If the substantial decrease in  $SP$  and increase in  $M$  are not roughly equivalent, then the overall result on  $R$  will be a substantial change in the direction of the variable that has a larger value. For example, if the absolute value of the decrease in  $SP$  is much greater than the absolute value of the increase in  $M$ , then  $R$  will substantially decrease.

Finally, there may be ways in which  $p$  can be changed from negative to positive, or its negative value can be reduced, by changing the way  $M$  is correlated to the action. To stay with the sex example, while there is some social discomfort associated with marriage “for money,” it does not approach the level of social disapprobation directed at prostitution. The  $p$  value is negative, but smaller. In other societies, perhaps in times holding less egalitarian ideals about marriage, there might have actually been a positive  $p$  value—as in a “good catch.” Similarly, professional performers or athletes may have been treated with less respect than amateurs a hundred years ago, but this has obviously changed quite dramatically. The same can be said for the *jalt* factor. One can imagine that free software-development communities would attach a negative social value to contributions of those who demand to be paid for their contributions. The same communities may have different feelings toward programmers who contributed for free but who later get large consulting contracts as a result of the experience and reputation they gained from their freely shared contributions.

This analysis suggests a series of likely conditions under which nonproprietary organizational approaches will be sustainable. First, there is the case of projects that are broken down into fine-grained modules, where market remuneration would likely be too costly to sustain, but where hedonic and social-psychological rewards can provide contributors with positive rewards. As I will explain in the following Subsection, fine-grained modularity is an important characteristic of the large-scale collaborations that form the basis of peer production. The analysis of motivations suggests

2002]

Coase's Penguin

that peer production likely will not be harnessed effectively using direct monetary incentives. Second, there are instances where the value of monetary return is small relative to the value of the hedonic and social-psychological rewards, particularly where the cultural construction of the social-psychological rewards places a high negative value on the direct association of monetary rewards with the activities. Teenagers and young adults with few economic commitments and a long time horizon for earning and saving, on the one hand, and high social recognition needs, on the other hand, are an obvious group fitting this description. Another group consists of individuals who have earnings sufficient to serve their present and expected tastes, but who have a strong taste for additional hedonic and social-psychological benefits that they could not obtain by extending their monetarily remunerated actions. Academics in general, and professional school academics in particular, are obvious instances of this group. Many of the volunteers for Internet-based projects who volunteer instead of watching television or reading a book likely fall into this category. Individuals whose present needs are met, but whose future expected needs require increased monetary returns, might participate if the social-psychological returns were not negatively correlated with future, indirect appropriation, such as reputation gains. This would effectively mean that they do add an *M* factor into their valuation of the rewards, but they do so in a way that does not negatively affect the value of *SP* for themselves or for other contributors to collaborative projects.

## 2. *Diverse Motivations and Large-Scale Collaborations*

The diversity of motivations allows large-scale collaborations to convert the motivation problem into a collaboration problem. In other words, the motivation problem is simple to resolve if the efforts of enough people can be pooled.

In a corollary to “Linus’s Law,”<sup>111</sup> one might say:

Given a sufficiently large number of contributions, direct monetary incentives necessary to bring about contributions are trivial.

The “sufficiently large” aspect of this observation requires some elaboration. “Sufficiently” refers to the fact that the number of people who need to collaborate to render the incentives problem trivial depends on the total cost or complexity of a project. The sustainability of any given project depends, therefore, not on the total cost but on how many individuals

---

111. Coined by Eric Raymond to capture one of the attributes of the approach that developed Linux: “Given enough eyeballs, all bugs are shallow.” RAYMOND, *supra* note 9, at 30.

contribute to it relative to the overall cost. If a project that requires thousands of person-hours can draw on the talents of fifteen- or thirty-thousand individuals instead of a few dozen or a few hundred, then the contribution of each, and hence the personal cost of participation that needs to be covered by diverse motivations, is quite low. Similarly, a project that requires ten or twenty person-hours can be provided with little heed to incentives if it can harness the distributed efforts of dozens of participants.

More generally, one can state:

Peer production is limited not by the total cost or complexity of a project, but by its modularity, granularity, and the cost of integration.

*Modularity* is a property of a project referring to the extent to which it can be broken down into smaller components, or modules, that can be independently and asynchronously produced before they are assembled into a whole. If modules are independent, individual contributors can choose what and when to contribute independently of each other, thereby maximizing their autonomy and flexibility to define the nature, extent, and timing of their participation in the project. Given the centrality of self-direction of human creative effort to the efficiencies of peer production, this characteristic is salient.

*Granularity* refers to the size of the modules, in terms of the time and effort that an agent must invest in producing them. The number of people who will likely participate in a project is inversely related to the size of the smallest-scale contribution necessary to produce a usable module. Usability may place a lower boundary on granularity either for technical or economic reasons, where at a minimum the cost of integrating a component into a larger modular project must be lower than the value that adding that component contributes to the project. But above that boundary, the granularity of the modules sets the smallest possible individual investment necessary to participate in a project. If this investment is sufficiently low, then incentives for producing that component of a modular project can be of trivial magnitude and many people can contribute. If the finest-grained contributions are relatively large and would require large investment of time and effort, the universe of potential contributors decreases. A predominant portion of the modules in a large-scale peer production project must therefore be relatively fine-grained for the project to be successful. The discussion in the preceding Subsection suggests that, given the relatively small independent value such fine-grained contributions will have and the transaction costs associated with remunerating each contribution monetarily, nonmonetary reward structures are likely to be more effective to motivate large-scale peer production efforts.

2002]

Coase's Penguin

Independent of the minimum granularity of a project, heterogeneity in the size of the modules may add to its efficiency. Heterogeneity allows contributors with diverse levels of motivation to collaborate by contributing modules of different sizes, whose production therefore requires different levels of motivation. Contributors may vary widely in their hedonic taste for creation, their social-psychological attitude toward participation, or in opportunities for indirect monetary appropriation (like the difference between IBM or Red Hat and individual volunteers in free software projects). A project that allows highly motivated contributors to carry a heavier load will be able to harness a diversely motivated human capital force more effectively than a project that can receive only standard-sized contributions.

#### B. *Integration: Problem and Opportunity*

The remaining obstacle to effective peer production is the problem of integration of the modules into a finished product. Integration includes two distinct components—first, a mechanism for providing quality control or integrity assurance to defend the project against incompetent or malicious contributions, and second, a mechanism for combining the contributed modules into a whole. It is here that the term “commons” that I use in describing the phenomenon as “commons-based peer production” gets its bite, denoting the centrality of the *absence of exclusion* as the organizing feature of this new mode of production and highlighting the potential pitfalls of such an absence for decentralized production. Observing commons-based peer production in the background of the commons literature, we see integration and the commons problem it represents solved in peer production efforts by a combination of four mechanisms: iterative peer production of the integration function itself, technical solutions embedded in the collaboration platform, norm-based social organization, and limited reintroduction of hierarchy or market to provide the integration function alone. In order for a project to be susceptible to sustainable peer production, the integration function must be either low-cost or itself sufficiently modular to be peer-produced in an iterative process.

Upon what kind of commons is it, then, that peer production of information relies? Commons are most importantly defined by two parameters.<sup>112</sup> The first parameter is whether use of the resource is common to everyone in the world or to a well-defined subset. The term “commons” is better reserved for the former, while the latter is better identified as a

---

112. The most extensive consideration of commons and the resolution of the collective action problems they pose is OSTROM, *supra* note 17.

common property regime (CPR)<sup>113</sup> or limited common property regime.<sup>114</sup> The second parameter is whether or not use of the resource by those whose use is privileged is regulated. Here one can state, following Carol Rose,<sup>115</sup> that resources in general can be subject to regimes ranging from total (and inefficiently delineated) exclusion—the phenomenon Michael Heller has called the anticommons<sup>116</sup>—through efficiently delineated property and otherwise regulated access, to completely open, unregulated access. The infamous “tragedy of the commons” is best reserved to refer only to the case of unregulated access commons, whether true commons or CPRs. Regulated commons need not be tragic at all, and indeed have been sustained and shown to be efficient in many cases.<sup>117</sup> The main difference is that CPRs are usually easier to monitor and regulate—using both formal law and social norms<sup>118</sup>—than true commons; hence, the latter may slip more often into the open access category even when they are formally regulated.

Elinor Ostrom also identified that one or both of two economic functions will be central to the potential failure or success of any given commons-based production system. The first is the question of provisioning, the second of allocation.<sup>119</sup> This may seem trivial, but it is important to keep the two problems separate, because if a particular resource is self-renewing when allocated properly, then institutions designed to assure provisioning would be irrelevant. Fishing and whaling are examples. In some cases, provisioning may be the primary issue.

---

113. See *id.*

114. Carol M. Rose, *The Several Futures of Property: Of Cyberspace and Folk Tales, Emission Trades and Ecosystems*, 83 MINN. L. REV. 129, 132 (1998).

115. Carol M. Rose, *Left Brain, Right Brain and History in the New Law and Economics of Property*, 79 OR. L. REV. 479, 480 (2000).

116. Michael A. Heller, *The Tragedy of the Anticommons: Property in the Transition from Marx to Markets*, 111 HARV. L. REV. 621 (1998). Although Frank Michelman originated the concept of the anticommons, see Frank I. Michelman, *Ethics, Economics, and the Law of Property*, in NOMOS XXIV: ETHICS, ECONOMICS, AND THE LAW 3, 6 (J. Roland Pennock & John W. Chapman eds., 1982), and Robert Ellickson coined the term “anticommons,” see Ellickson, *supra* note 70, at 1322 n.22, I refer to Heller, because the concept took off with Heller’s use. See PERSPECTIVES ON PROPERTY LAW 167 (Robert C. Ellickson et al. eds., 2002) (describing the history of the term “anticommons”).

117. For the most comprehensive survey, see OSTROM, *supra* note 17. Another seminal study was JAMES M. ACHESON, THE LOBSTER GANGS OF MAINE (1988). A brief intellectual history of the study of common resource pools and common property regimes can be found in Charlotte Hess & Elinor Ostrom, Artifacts, Facilities, and Content: Information as a Common-Pool Resource, Paper Presented at the Conference on the Public Domain, Duke Law School, Nov. 9-11, 2001, at <http://www.law.duke.edu/pd/papers/ostromhes.pdf>.

118. The particular focus on social norms rather than on formal regulation as central to the sustainability of common resource pool management solutions that are not based on property is Ellickson’s. See ROBERT C. ELICKSON, ORDER WITHOUT LAW: HOW NEIGHBORS SETTLE DISPUTES (1991); Ellickson, *supra* note 70.

119. “Provisioning” refers to efforts aimed at producing a particular good that would not otherwise exist. “Allocating” refers to decisions about how a good that exists, but is scarce relative to demand for it, will be used most efficiently.

2002]

## Coase's Penguin

Ostrom describes various water districts that operate as common property regimes that illustrate well the differences between situations where allocation of a relatively stable (but scarce) water flow is the problem, on one hand, and those where provisioning of a dam (which will result in abundant water supply, relative to demand) is the difficult task.<sup>120</sup> Obviously, some commons require both operations.

Information production entails only a provisioning problem. Because information is nonrival, once it is produced no allocation problem exists. Moreover, commons-based provisioning of information in a ubiquitously networked environment may present a more tractable problem than provisioning of physical matter, and shirking or free-riding may not lead quite as directly to nonproduction. This is so for three reasons. First, the modularity of the projects allows redundant provisioning of “dropped” components to overcome occasional defections without threatening the whole. Second, a ubiquitously networked environment substantially increases the size of the pool of contributors. At first glance, this should undermine peer production, because the likelihood of free-riding generally increases as the size of the pool increases and the probability of social-norms-based prevention of free-riding declines.<sup>121</sup> But as the size of the pool increases, the project can tolerate increasing levels of free-riding, as long as the absolute number of contributors responding to some mix of motivations remains sufficiently large such that the aggregation of the efforts of those who do contribute, each at a level no higher than his or her level of motivation dictates, will be adequate to produce the good. As long as free-riders do not affirmatively undermine production but simply do not contribute, the willingness of contributors to contribute should depend on their perception of the likelihood of success given the number of contributors, not on the total number of users. Indeed, for contributors who seek indirect appropriation through means enhanced by widespread use of the joint product—like reputation or service contracts—a high degree of use of the end product, even by “free-riders” who did not contribute to providing it, increases the expected payoff.<sup>122</sup> Third, the public goods nature of the product means that free-riding does not affect the capacity of contributors to gain full use of their joint product and does not degrade their

---

120. OSTRUM, *supra* note 17, at 69-88.

121. On the relationship between how small and closely knit a group is and its capacity to use social norms to regulate behavior, see ELICKSON, *supra* note 118. On the importance of social norms in regulating behavior in cyberspace, see LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE (1999).

122. This attribute causes Steven Weber to describe free software production not only as “nonrival,” but as “anti-rival”—by which he means that increasing returns to widespread use mean that consumption by many not only does not reduce the value of a good, as in nonrival goods, but actually enhances it. See WEBER, *supra* note 8, at 28-29.

utility from it. This permits contributors who contribute in expectation of the use value of the good to contribute without concern for free-riding.

There are, however, types of defection that are likely to undermine provisioning by adversely affecting either (1) motivation to participate or (2) the efficacy of participation. The first type covers actions that reduce the value of participation, be it the intrinsic hedonic or social-psychological components or the expected longer-term extrinsic values, that is, the monetary rewards to reputation, human capital, etc. The second type relates mostly to potential failures of integration, such as poor quality-control mechanisms.

### 1. *Threats to Motivation*

There are two kinds of actions that could reduce the intrinsic benefits of participation. First is the possibility that a behavior will affect the contributors' valuation of the intrinsic value of participation. Two primary sources of negative effect seem likely. The first is a failure of integration, so that the act of individual provisioning is seen as being wasted, rather than adding some value to the world. This assumes that contributors have a taste for contributing to a successful joint project. Where this is not the case—if integration is not a component of the intrinsic value of participation—then failure to integrate would not be significant. The World Wide Web is an example where it is quite possible that putting up a website on a topic one cares about is sufficiently intrinsically valuable to the author, even without the sense of adding to the great library of the Web, that integration is irrelevant to the considerations of many contributors.

The more important potential “defection” from commons-based peer production is unilateral<sup>123</sup> appropriation. Unilateral appropriation could, but need not, take the form of commercialization of the common efforts for private benefit. More generally, appropriation could be any act where an individual contributor tries to make the common project reflect his or her values too much, thereby alienating other participants from the product of their joint effort. The common storytelling enterprise called LambdaMOO encountered well-described crises with individuals who behaved in antisocial ways, taking control over the production process to make the joint product serve their own goals—like forcing female characters in the story to “have sex” that they did not want to have.<sup>124</sup> In LambdaMOO, the participants set up a social structure for clearing common political will in

---

123. Unilateral is opposed to collective, as in the conversion of some aspect of the commons to a common property regime where, for example, high quality or consistent contribution to the commons could become a criterion for membership.

124. LESSIG, *supra* note 121, at 74-75.

2002]

Coase's Penguin

response to this form of appropriation.<sup>125</sup> In the examples I have described in this Article, the explicit adherence to a norm of objectivity in Wikipedia<sup>126</sup> and similar references in Kuro5hin to the norm of high-quality writing are clear examples of efforts to use social norms to regulate this type of defection by substantive, rather than commercial, appropriation. Similarly, some of the software-based constraints on moderation and commenting on Slashdot and other sites serve to prevent anyone from taking too large a role in shaping the direction of the common enterprise in a way that would reduce the perceived benefits of participation to many others. For example, limiting moderators to moderating no more than five comments in any three-day period or using troll filters to prevent users from posting too often are technical constraints that do not permit anyone to appropriate the common enterprise called Slashdot.

Another form of appropriation that could affect valuation of participation is simple commercialization for private gain. The primary concern is that commercialization by some participants or even by nonparticipants will create a sucker's reward aspect to participation. This is the effect I introduced into the abstract statement of diverse motivations as the *jalt* factor—the effect of monetary rewards for others on the perceived value of participation. One example of such an effect may have occurred when the early discussion moderators on AOL boards—volunteers all—left when they began to realize that their contributions were effectively going to increase the value of the company. There is, however, an immensely important counterexample—to wit, the apparent imperviousness of free software production, our paradigm case, to this effect. Some contributors have made billions, while some of the leaders of major projects have earned nothing but honor.<sup>127</sup> Query, though, whether the pattern would have held if the primary leader of a project, such as Linus Torvalds, rather than people less central to the Linux kernel development process, had made money explicitly by selling the GNU/Linux operating system as a product. It is, in any event, not implausible to imagine that individuals would be more willing to contribute their time and effort to NASA or a nonprofit enterprise than to a debugging site set up by Microsoft. Whether the *jalt* effect exists, how strong it is, and what are the characteristics of instances when it is or is not important is a valuable area for empirical research.

In addition to the intrinsic value of participation, an important component of motivation also relies on the use value of the joint project and

---

125. See Julian Dibbell, *A Rape in Cyberspace or How an Evil Clown, a Haitian Trickster Spirit, Two Wizards, and a Cast of Dozens Turned a Database into a Society*, VILLAGE VOICE, Dec. 21, 1993, at 36, at <ftp://ftp.lambda.moo.mud.org/pub/MOO/papers/VillageVoice.txt>.

126. See *supra* text accompanying note 35.

127. In this, too, peer production is similar to academic production, where scientists see their basic research used, very often by others, as the basis for great wealth in which they do not share.

on indirect appropriation based on continued access to the joint product—service contracts and human capital for instance. For such projects, defection may again take the form of appropriation, in this case by exclusion of contributors from the use value of the end product. (Why academics, for example, are willing to accept the bizarre system in which they contribute to peer review journals for free, sometimes even paying a publication fee, and then have their institutions buy this work back from the printers at exorbitant rates remains a mystery.) In free software, the risk of defection through this kind of appropriation is deemed a central threat to the viability of the enterprise, and the GNU GPL is designed precisely to prevent one person from taking from the commons, appropriating the software, and excluding others from it.<sup>128</sup> This type of defection looks like an allocation problem—one person is taking more than his or her fair share. But again, this is true only in a metaphoric sense. The good is still intrinsically a public good and is physically available to be used by everyone. Law (intellectual property) may create this “allocation problem” in a misguided attempt to solve a perceived provisioning problem, but the real problem is the effect on motivation to provision, not an actual scarcity that requires better allocation. The risk of this kind of unilateral appropriation lowers the expected value contributors can capture from their contribution, and hence lowers motivation to participate and provide the good.

## 2. *Provisioning Integration*

Another potential problem that commons-based peer production faces is provisioning of the integration function itself. It is important to understand that integration requires some process for assuring the quality of individual contributions. This could take the form of (1) hierarchically managed review, as in the Linux kernel or Apache development processes; (2) peer review, as in the process for moderating Slashdot comments; (3) norms-based social organization, as in Wikipedia’s objectivity norm; or (4) aggregation and averaging of redundant contributions, as in the Mars Clickworkers project. Academic peer production of science is traditionally some combination of the first three, although the Los Alamos Archive<sup>129</sup> and the Varmus proposal for changing the model of publication in the

---

128. Free Software Found., *supra* note 19. Section 2(b) limits the license to modify software distributed under the GPL such that the licensee “must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.” *Id.*

129. ArXiv.org, ArXiv.org E-Print Archive, *at* <http://www.arxiv.org> (last visited Sept. 25, 2002).

2002]

## Coase's Penguin

health and biomedical sciences<sup>130</sup> to free online publication coupled with post-publication peer commentary as a check on quality would tend to push the process further toward pure peer review and norms-based enforcement of the core values of completeness and accuracy, as well as attribution and respect for priority.

The first thing to see from the discussion of threats to motivation is that provisioning integration by permitting the integrator to be the residual owner (in effect, to “hire” the contributors and act as the entrepreneur) presents substantial problems for the motivation to contribute in a peer-based production model. Appropriation may so affect motivation to participate that the residual owner will have to resort to market- and hierarchy-based organization of the whole production effort. Second, property rights in information are always in some measure inefficient.<sup>131</sup> Creating full property rights in any single agent whose contribution is only a fraction of the overall investment in the product is even less justifiable than doing so for a person who pays all of the production costs. Third, integration is quite possibly a low-cost activity, particularly with the introduction of software-based management of the communications and to some extent the integration of effort. To the extent that this is so, even though integration may require some hierarchy or some market-based provisioning, it is a function that can nonetheless be sustained on low returns by volunteers, like those who run integration of code into the Linux kernel or by publicly funded actors, as in the case of NASA Clickworkers. It can also be sustained by firms, like the Open Source Development Network that supports Slashdot, that rely on business models that do not depend on intellectual property rights.

The cost of integration—and hence the extent to which it is a limit on the prevalence of peer production—can be substantially reduced by automation and the introduction of an iterative process of peer production of integration itself. First, integration could be a relatively automated process for some products. The use of automated collation of markings and averaging out of deviations by NASA Clickworkers is an example, as are many of the attributes of Slashdot or Kuro5hin. Second, the integration function itself can be peer-produced. Again with Slashdot, the software that provides important integration functions is itself an open source project—in other words, peer-produced. The peer review of the peer reviewers—the moderators—is also distributed, in that ninety percent of registered users can review the moderators, who in turn review the contributors. As peer production is iteratively introduced to solve a greater portion of the

---

130. Harold Varmus, E-BIOMED: A Proposal for Electronic Publications in the Biomedical Sciences (1999), at <http://www.nih.gov/about/director/pubmedcentral/ebiomedarch.htm>.

131. See *supra* note 74 and accompanying text.

integration function, the residual investment in integration that might require some other centralized provisioning becomes a progressively smaller investment, one that is capable of being carried on by volunteers or by firms that need not appropriate anything approaching the full value of the product.<sup>132</sup>

Moreover, integration could provide an opportunity for cooperative monetary appropriation, not only, or even primarily, integration into a general product but integration as a specific customization for specific users.<sup>133</sup> There are no models for such cooperative appropriation on a large scale yet, but the idea is that many peers will be admitted to something that is more akin to a common property regime or partnership than a commons, probably on the basis of reputation for contributing to the commons, and these groups would develop a system for receiving and disseminating service/customization projects (if it is a software project) or other information-production projects. This would not necessarily work for all information production, but it could work in some. The idea is that the indirect appropriation itself would be organized on a peer model so that reputation would lead not to being hired as an employee by a hierarchical firm, but would instead be an initiation into a cooperative, managed and “owned” by its participants. Just as in the case of Slashdot, some mechanism for assuring quality of work in the products would be necessary, but it would be achievable on a distributed model, rather than a hierarchical model, with some tracking of individual contribution to any given project (or some other mechanism for distribution of revenues). The idea here would be to provide a peer-based model for allowing contributors to share the benefits of large-scale service projects, rather than relegating them to individual indirect appropriation.

To conclude, whether or not a peer production project will be able to resolve the integration problem is a central limiting factor on the viability of peer production to provision any given information goods. Approaches to integration include technology, as with the software running Slashdot or the Clickworkers project; iterative peer production, such as the moderation and meta-moderation on Slashdot; social norms, as with Wikipedia’s or Kuro5hin; and market or hierarchical mechanisms that integrate the project without appropriating the joint product, as is the case in the Linux kernel development community.

---

132. Boyle focuses on this characteristic as the most interesting and potentially important solution. See Boyle, *supra* note 24, at 13.

133. I owe the idea of cooperative monetary appropriation to an enormously productive conversation with David Johnson. It was his idea that the peer production model can be combined with the producers’ cooperative model to provide a mechanism of appropriation that would give contributors to peer production processes a more direct mechanism for keeping body and soul together while contributing, rather than simply waiting for reputation gains to be translated into a contract with a company.

2002]

Coase's Penguin

## CONCLUSION

In this Article, I suggest that peer production of information is a phenomenon with much broader economic implications for information production than thinking of free software alone would suggest. I describe commons-based peer production enterprises occurring throughout the value chain of information production on the Internet, from content production, through relevance and accreditation, to distribution. I then explain that peer production has a systematic advantage over markets and firms in matching the best available human capital to the best available information inputs in order to create information products.

Peer production of information is emerging because both the declining price of physical capital involved in information production and the declining price of communications lower the cost of peer production and make human capital the primary economic good involved. These trends both lower the cost of coordination and increase the importance of peer production's relative advantage—identifying the best available human capital in highly refined increments and allocating it to projects. If true, this phenomenon has a number of implications both for firms seeking to structure a business model for the Internet and for governments seeking to capitalize on the Internet to become more innovative and productive.

For academics, peer production provides a rich area for new research. Peer production, like the Internet, is just emerging. While there are some studies of peer-produced software, there is little by way of systematic research into peer production processes more generally. There is much room for theoretical work on why they are successful, as well as potential pitfalls and the solutions that, in principle and in practice, can be adopted in response to those pitfalls. The role of norms, the role of technology, and the interaction between volunteerism and economic gain in shaping the motivation and organization of peer production are also important areas of research, in particular for the study of how peer groups cluster around projects. Qualitative and quantitative studies of the importance of peer production in the overall information economy, in particular the Internet-based information economy, would provide a better picture of just how central or peripheral a phenomenon this is.

For firms, the emergence of peer production may require a more aggressive move from information product-based business models to information-embedding material products and service-based business models. Businesses could, following IBM or Red Hat in open source software, focus their “production” investment on providing opportunities for peer production, aiding in that production, and performing some of the integration functions. Firms that adopt this model, however, will not be able to count on appropriating the end product directly, because the threat of

appropriation will largely dissipate motivations for participation. Indeed, the capacity of a firm to commit credibly *not* to appropriate the joint project will be crucial to its success in building a successful business model involving a peer production process. This commitment would require specific licenses that secure access to the work over time to everyone, including contributors. It would also require a business model that depends on indirect appropriation of the benefits of the product.<sup>134</sup> Selling products or services for which availability of the peer-produced product increases demand could do this, as in the case of IBM servers that run the GNU/Linux operating system and Apache server software. Conversely, firms that benefit on the supply side from access to certain types of information can capitalize on peer production processes to provide that input cheaply and efficiently, while gaining the firm-specific human capital to optimize their product to fit the information. Again, IBM's investment in engineers who participate in writing open source software releases it from reliance on proprietary software owned by other firms, thereby creating supply-side economies to its support of peer production of software. Similarly, NASA's utilization of peer production reduces its costs of mapping Mars craters, and Google's use of links provided by websites as votes for relevance integrates distributed relevance judgments as input into its own commercial product. Another option is sale of the tools of peer production itself. For example, the popularity of software and access to massive multiplayer online games like Ultima Online or Everquest are instances of a growing industry in the tools for peer production of escapist storytelling.<sup>135</sup>

For regulators, the implications are quite significant. In particular, the current heavy focus on strengthening intellectual property rights is exactly the wrong approach to increasing growth through innovation and information production if having a robust peer production sector is important to an economy's capacity to tap its human capital efficiently. Strong intellectual property rights, in particular rights to control creative utilization of existing information, harm peer production by raising the cost of access to existing information resources as input. This barrier limits the capacity of the hundreds of thousands of potential contributors to consider what could be done with a given input and to apply themselves to it without violating the rights of the owner of the information input. This does not mean that intellectual property rights are entirely bad. But we have known

---

134. For a general mapping of indirect appropriation mechanisms, see Benkler, *supra* note 76.

135. See Bob Tedeschi, *E-Commerce Report: The Computer Game Industry Seeks To Bridge an Online Gap Between Geeks and the Mainstream*, N.Y. TIMES, Dec. 31, 2001, at C5. A somewhat optimistic report estimates that this industry will pull in some \$1.3 billion by 2006. See Tamsin McMahon, *Gaming Platforms Set for Explosive Growth*, EUROPEMEDIA, July 3, 2002, at <http://www.europemedia.net/showness.asp?ArticleID=11326>.

2002]

Coase's Penguin

for decades that intellectual property entails systematic inefficiencies as a solution to the problem of private provisioning of the public good called information. The emergence of commons-based peer production adds a new source of inefficiency.

The strength of peer production is in matching human capital to information inputs to produce new information goods. Strong intellectual property rights inefficiently shrink the universe of existing information inputs that can be subjected to this process. Instead, owned inputs will be limited to human capital with which the owner of the input has a contractual—usually employment—relationship. Moreover, the entire universe of peer-produced information gains no benefit from strong intellectual property rights. Since the core of commons-based peer production entails provisioning without direct appropriation and since indirect appropriation—intrinsic or extrinsic—does not rely on control of the information but on its widest possible availability, intellectual property offers no gain, only loss, to peer production. While it is true that free software currently uses copyright-based licensing to prevent certain kinds of defection from peer production processes, that strategy is needed only as a form of institutional jiu-jitsu to defend from intellectual property.<sup>136</sup> A complete absence of property in the software domain would be at least as congenial to free software development as the condition where property exists, but copyright permits free software projects to use licensing to defend themselves from defection. The same protection from defection might be provided by other means as well, such as creating simple public mechanisms for contributing one's work in a way that makes it unsusceptible to downstream appropriation—a conservancy of sorts. Regulators concerned with fostering innovation may better direct their efforts toward providing the institutional tools that would help thousands of people to collaborate without appropriating their joint product, making the information they produce freely available rather than spending their efforts to increase the scope and sophistication of the mechanisms for private appropriation of this public good as they now do.

That we cannot fully understand a phenomenon does not mean that it does not exist. That a seemingly growing phenomenon refuses to fit our longstanding perceptions of how people behave and how economic growth occurs counsels closer attention, not studied indifference and ignorance. Commons-based peer production presents a fascinating phenomenon that could allow us to tap substantially underutilized reserves of human creative effort. It is of central importance that we not squelch peer production, but that we create the institutional conditions needed for it to flourish.

---

136. *But see* McGowan, *supra* note 13, at 287-88 (arguing that the right to exclude will always be necessary to prevent opportunistic appropriation of open source code).